

WO 00/13091

PCT/US98/24943

INTELLIGENT NETWORK INTERFACE DEVICE
AND SYSTEM FOR ACCELERATING COMMUNICATION

Technical Field

5 The present invention relates generally to computer or other networks, and more particularly to processing of information communicated between hosts such as computers connected to a network.

Background

10 The advantages of network computing are increasingly evident. The convenience and efficiency of providing information, communication or computational power to individuals at their personal computer or other end user devices has led to rapid growth of such network computing, including internet as well as intranet devices and applications.

 As is well known, most network computer communication is accomplished with the aid of a layered software architecture for moving information between host computers
15 connected to the network. The layers help to segregate information into manageable segments, the general functions of each layer often based on an international standard called Open Systems Interconnection (OSI). OSI sets forth seven processing layers through which information may pass when received by a host in order to be presentable to an end user. Similarly, transmission of information from a host to the network may pass through those
20 seven processing layers in reverse order. Each step of processing and service by a layer may include copying the processed information. Another reference model that is widely implemented, called TCP/IP (TCP stands for transport control protocol, while IP denotes internet protocol) essentially employs five of the seven layers of OSI.

 Networks may include, for instance, a high-speed bus such as an Ethernet connection
25 or an internet connection between disparate local area networks (LANs), each of which includes multiple hosts, or any of a variety of other known means for data transfer between hosts. According to the OSI standard, physical layers are connected to the network at respective hosts, the physical layers providing transmission and receipt of raw data bits via the network. A data link layer is serviced by the physical layer of each host, the data link
30 layers providing frame division and error correction to the data received from the physical layers, as well as processing acknowledgment frames sent by the receiving host. A network layer of each host is serviced by respective data link layers, the network layers primarily controlling size and coordination of subnets of packets of data.

WO 00/13091

PCT/US98/24943

A transport layer is serviced by each network layer and a session layer is serviced by each transport layer within each host. Transport layers accept data from their respective session layers and split the data into smaller units for transmission to the other host's transport layer, which concatenates the data for presentation to respective presentation layers.

- 5 Session layers allow for enhanced communication control between the hosts. Presentation layers are serviced by their respective session layers, the presentation layers translating between data semantics and syntax which may be peculiar to each host and standardized structures of data representation. Compression and/or encryption of data may also be accomplished at the presentation level. Application layers are serviced by respective
- 10 presentation layers, the application layers translating between programs particular to individual hosts and standardized programs for presentation to either an application or an end user. The TCP/IP standard includes the lower four layers and application layers, but integrates the functions of session layers and presentation layers into adjacent layers. Generally speaking, application, presentation and session layers are defined as upper layers,
- 15 while transport, network and data link layers are defined as lower layers.

- The rules and conventions for each layer are called the protocol of that layer, and since the protocols and general functions of each layer are roughly equivalent in various hosts, it is useful to think of communication occurring directly between identical layers of different hosts, even though these peer layers do not directly communicate without
- 20 information transferring sequentially through each layer below. Each lower layer performs a service for the layer immediately above it to help with processing the communicated information. Each layer saves the information for processing and service to the next layer. Due to the multiplicity of hardware and software architectures, devices and programs commonly employed, each layer is necessary to insure that the data can make it to the
- 25 intended destination in the appropriate form, regardless of variations in hardware and software that may intervene.

- In preparing data for transmission from a first to a second host, some control data is added at each layer of the first host regarding the protocol of that layer, the control data being indistinguishable from the original (payload) data for all lower layers of that host. Thus an
- 30 application layer attaches an application header to the payload data and sends the combined data to the presentation layer of the sending host, which receives the combined data, operates on it and adds a presentation header to the data, resulting in another combined data packet. The data resulting from combination of payload data, application header and presentation header is then passed to the session layer, which performs required operations including

WO 00/13091

PCT/US98/24943

attaching a session header to the data and presenting the resulting combination of data to the transport layer. This process continues as the information moves to lower layers, with a transport header, network header and data link header and trailer attached to the data at each of those layers, with each step typically including data moving and copying, before sending
5 the data as bit packets over the network to the second host.

The receiving host generally performs the converse of the above-described process, beginning with receiving the bits from the network, as headers are removed and data processed in order from the lowest (physical) layer to the highest (application) layer before transmission to a destination of the receiving host. Each layer of the receiving host
10 recognizes and manipulates only the headers associated with that layer, since to that layer the higher layer control data is included with and indistinguishable from the payload data. Multiple interrupts, valuable central processing unit (CPU) processing time and repeated data copies may also be necessary for the receiving host to place the data in an appropriate form at its intended destination.

15 The above description of layered protocol processing is simplified, as college-level textbooks devoted primarily to this subject are available, such as Computer Networks, Third Edition (1996) by Andrew S. Tanenbaum, which is incorporated herein by reference. As defined in that book, a computer network is an interconnected collection of autonomous computers, such as internet and intranet devices, including local area networks (LANs), wide
20 area networks (WANs), asynchronous transfer mode (ATM), ring or token ring, wired, wireless, satellite or other means for providing communication capability between separate processors. A computer is defined herein to include a device having both logic and memory functions for processing data, while computers or hosts connected to a network are said to be heterogeneous if they function according to different operating devices or communicate via
25 different architectures.

As networks grow increasingly popular and the information communicated thereby becomes increasingly complex and copious, the need for such protocol processing has increased. It is estimated that a large fraction of the processing power of a host CPU may be devoted to controlling protocol processes, diminishing the ability of that CPU to perform
30 other tasks. Network interface cards have been developed to help with the lowest layers, such as the physical and data link layers. It is also possible to increase protocol processing speed by simply adding more processing power or CPUs according to conventional arrangements. This solution, however, is both awkward and expensive. But the complexities presented by various networks, protocols, architectures, operating devices and applications

WO 00/13091

PCT/US98/24943

generally require extensive processing to afford communication capability between various network hosts.

Disclosure of the Invention

The current invention provides a device for processing network communication that
5 greatly increases the speed of that processing and the efficiency of transferring data being communicated. The invention has been achieved by questioning the long-standing practice of performing multilayered protocol processing on a general-purpose processor. The protocol processing method and architecture that results effectively collapses the layers of a
10 connection-based, layered architecture such as TCP/IP into a single wider layer which is able to send network data more directly to and from a desired location or buffer on a host. This accelerated processing is provided to a host for both transmitting and receiving data, and so improves performance whether one or both hosts involved in an exchange of information have such a feature.

The accelerated processing includes employing representative control instructions for
15 a given message that allow data from the message to be processed via a fast-path which accesses message data directly at its source or delivers it directly to its intended destination. This fast-path bypasses conventional protocol processing of headers that accompany the data. The fast-path employs a specialized microprocessor designed for processing network communication, avoiding the delays and pitfalls of conventional software layer processing,
20 such as repeated copying and interrupts to the CPU. In effect, the fast-path replaces the states that are traditionally found in several layers of a conventional network stack with a single state machine encompassing all those layers, in contrast to conventional rules that require rigorous differentiation and separation of protocol layers. The host retains a sequential protocol processing stack which can be employed for setting up a fast-path connection or
25 processing message exceptions. The specialized microprocessor and the host intelligently choose whether a given message or portion of a message is processed by the microprocessor or the host stack.

Brief Description of the Drawings

FIG. 1 is a plan view diagram of a device of the present invention, including a host
30 computer having a communication-processing device for accelerating network communication.

FIG. 2 is a diagram of information flow for the host of FIG. 1 in processing network communication, including a fast-path, a slow-path and a transfer of connection context between the fast and slow-paths.

WO 00/13091

PCT/US98/24943

FIG. 3 is a flow chart of message receiving according to the present invention.

FIG. 4A is a diagram of information flow for the host of FIG. 1 receiving a message packet processed by the slow-path.

FIG. 4B is a diagram of information flow for the host of FIG. 1 receiving an initial
5 message packet processed by the fast-path.

FIG. 4C is a diagram of information flow for the host of FIG. 4B receiving a subsequent message packet processed by the fast-path.

FIG. 4D is a diagram of information flow for the host of FIG. 4C receiving a message packet having an error that causes processing to revert to the slow-path.

FIG. 5 is a diagram of information flow for the host of FIG. 1 transmitting a message
10 by either the fast or slow-paths.

FIG. 6 is a diagram of information flow for a first embodiment of an intelligent network interface card (INIC) associated with a client having a TCP/IP processing stack.

FIG. 7 is a diagram of hardware logic for the INIC embodiment shown in FIG. 6,
15 including a packet control sequencer and a fly-by sequencer.

FIG. 8 is a diagram of the fly-by sequencer of FIG. 7 for analyzing header bytes as they are received by the INIC.

FIG. 9 is a diagram of information flow for a second embodiment of an INIC associated with a server having a TCP/IP processing stack.

FIG. 10 is a diagram of a command driver installed in the host of FIG. 9 for creating
20 and controlling a communication control block for the fast-path.

FIG. 11 is a diagram of the TCP/IP stack and command driver of FIG. 10 configured for NetBios communications.

FIG. 12 is a diagram of a communication exchange between the client of FIG. 6 and
25 the server of FIG. 9.

FIG. 13 is a diagram of hardware functions included in the INIC of FIG. 9.

FIG. 14 is a diagram of a trio of pipelined microprocessors included in the INIC of FIG. 13, including three phases with a processor in each phase.

FIG. 15A is a diagram of a first phase of the pipelined microprocessor of FIG. 14.

FIG. 15B is a diagram of a second phase of the pipelined microprocessor of FIG. 14.
30

FIG. 15C is a diagram of a third phase of the pipelined microprocessor of FIG. 14.

FIG. 16 is a diagram of a plurality of queue storage units that interact with the microprocessor of FIG. 14 and include SRAM and DRAM.

FIG. 17 is a diagram of a set of status registers for the queues storage units of FIG. 16.

WO 00/13091

PCT/US98/24943

FIG. 18 is a diagram of a queue manager, which interacts, with the queue storage units and status registers of FIG. 16 and FIG. 17.

FIGs. 19A-D are diagrams of various stages of a least-recently-used register that is employed for allocating cache memory.

5 FIG. 20 is a diagram of the devices used to operate the least-recently-used register of FIGs. 19A-D.

Best Mode for Carrying Out the Invention

FIG. 1 shows a host 20 of the present invention connected by a network 25 to a remote host 22. The increase in processing speed achieved by the present invention can be provided with an intelligent network interface card (INIC) that is easily and affordably added to an existing host, or with a communication processing device (CPD) that is integrated into a host, in either case freeing the host CPU from most protocol processing and allowing improvements in other tasks performed by that CPU. The host 20 in a first embodiment contains a CPU 28 and a CPD 30 connected by a PCI bus 33. The CPD 30 includes a microprocessor designed for processing communication data and memory buffers controlled by a direct memory access (DMA) unit. Also connected to the PCI bus 33 is a storage device 35, such as a semiconductor memory or disk drive, along with any related controls.

Referring additionally to FIG. 2, the host CPU 28 controls a protocol processing stack 44 housed in storage 35, the stack including a data link layer 36, network layer 38, transport layer 40, upper layer 46 and an upper layer interface 42. The upper layer 46 may represent a session, presentation and/or application layer, depending upon the particular protocol being employed and message communicated. The upper layer interface 42, along with the CPU 28 and any related controls can send or retrieve a file to or from the upper layer 46 or storage 35, as shown by arrow 48. A connection context 50 has been created, as will be explained below, the context summarizing various features of the connection, such as protocol type and source and destination addresses for each protocol layer. The context may be passed between an interface for the session layer 42 and the CPD 30, as shown by arrows 52 and 54, and stored as a communication control block (CCB) at either CPD 30 or storage 35.

When the CPD 30 holds a CCB defining a particular connection, data received by the CPD from the network and pertaining to the connection is referenced to that CCB and can then be sent directly to storage 35 according to a fast-path 58, bypassing sequential protocol processing by the data link 36, network 38 and transport 40 layers. Transmitting a message, such as sending a file from storage 35 to remote host 22, can also occur via the fast-path 58, in which case the context for the file data is added by the CPD 30 referencing a CCB, rather

WO 00/13091

PCT/US98/24943

than by sequentially adding headers during processing by the transport 40, network 38 and data link 36 layers. The DMA controllers of the CPD 30 perform these transfers between CPD and storage 35.

5 The CPD 30 collapses multiple protocol stacks each having possible separate states into a single state machine for fast-path processing. As a result, exception conditions may occur that are not provided for in the single state machine, primarily because such conditions occur infrequently and to deal with them on the CPD would provide little or no performance benefit to the host. Such exceptions can be CPD 30 or CPU 28 initiated. An advantage of the invention includes the manner in which unexpected situations that occur on a fast-path
10 CCB are handled. The CPD 30 deals with these rare situations by passing back or flushing to the host protocol stack 44 the CCB and any associated message frames involved, via a control negotiation. The exception condition is then processed in a conventional manner by the host protocol stack 44. At some later time, usually directly after the handling of the exception condition has completed and fast-path processing can resume, the host stack 44 hands the
15 CCB back to the CPD.

This fallback capability enables the performance-impacting functions of the host protocols to be handled by the CPD network microprocessor, while the exceptions are dealt with by the host stacks, the exceptions being so rare as to negligibly effect overall performance. The custom designed network microprocessor can have independent
20 processors for transmitting and receiving network information, and further processors for assisting and queuing. A preferred microprocessor embodiment includes a pipelined trio of receive, transmit and utility processors. DMA controllers are integrated into the implementation and work in close concert with the network microprocessor to quickly move data between buffers adjacent to the controllers and other locations such as long term storage.
25 Providing buffers logically adjacent to the DMA controllers avoids unnecessary loads on the PCI bus.

FIG. 3 diagrams the general flow of messages received according to the current invention. A large TCP/IP message such as a file transfer may be received by the host from the network in a number of separate, approximately 64 KB transfers, each of which may be
30 split into many, approximately 1.5 KB frames or packets for transmission over a network. Novell NetWare protocol suites running Sequenced Packet Exchange Protocol (SPX) or NetWare Core Protocol (NCP) over Internetwork Packet Exchange (IPX) work in a similar fashion. Another form of data communication which can be handled by the fast-path is Transaction TCP (hereinafter T/TCP or TTCP), a version of TCP which initiates a connection

WO 00/13091

PCT/US98/24943

with an initial transaction request after which a reply containing data may be sent according to the connection, rather than initiating a connection via a several-message initialization dialogue and then transferring data with later messages. In any of the transfers typified by these protocols, each packet conventionally includes a portion of the data being transferred, as well as headers for each of the protocol layers and markers for positioning the packet relative to the rest of the packets of this message.

When a message packet or frame is received 47 from a network by the CPD, it is first validated by a hardware assist. This includes determining the protocol types of the various layers, verifying relevant checksums, and summarizing 57 these findings into a status word or words. Included in these words is an indication whether or not the frame is a candidate for fast-path data flow. Selection 59 of fast-path candidates is based on whether the host may benefit from this message connection being handled by the CPD, which includes determining whether the packet has header bytes indicating particular protocols, such as TCP/IP or SPX/IPX for example. The small percent of frames that are not fast-path candidates are sent 61 to the host protocol stacks for slow-path protocol processing. Subsequent network microprocessor work with each fast-path candidate determines whether a fast-path connection such as a TCP or SPX CCB is already extant for that candidate, or whether that candidate may be used to set up a new fast-path connection, such as for a TTCP/IP transaction. The validation provided by the CPD provides acceleration whether a frame is processed by the fast-path or a slow-path, as only error free, validated frames are processed by the host CPU even for the slow-path processing.

All received message frames which have been determined by the CPD hardware assist to be fast-path candidates are examined 53 by the network microprocessor or INIC comparator circuits to determine whether they match a CCB held by the CPD. Upon confirming such a match, the CPD removes lower layer headers and sends 69 the remaining application data from the frame directly into its final destination in the host using direct memory access (DMA) units of the CPD. This operation may occur immediately upon receipt of a message packet, for example when a TCP connection already exists and destination buffers have been negotiated, or it may first be necessary to process an initial header to acquire a new set of final destination addresses for this transfer. In this latter case, the CPD will queue subsequent message packets while waiting for the destination address, and then DMA the queued application data to that destination.

A fast-path candidate that does not match a CCB may be used to set up a new fast-path connection, by sending 65 the frame to the host for sequential protocol processing. In

WO 00/13091

PCT/US98/24943

this case, the host uses this frame to create 51 a CCB, which is then passed to the CPD to control subsequent frames on that connection. The CCB, which is cached 67 in the CPD, includes control and state information pertinent to all protocols that would have been processed had conventional software layer processing been employed. The CCB also
 5 contains storage space for per-transfer information used to facilitate moving application-level data contained within subsequent related message packets directly to a host application in a form available for immediate usage. The CPD takes command of connection processing upon receiving a CCB for that connection from the host.

As shown more specifically in FIG. 4A, when a message packet is received from the
 10 remote host 22 via network 25, the packet enters hardware receive logic 32 of the CPD 30, which checksums headers and data, and parses the headers, creating a word or words which identify the message packet and status, storing the headers, data and word temporarily in memory 60. As well as validating the packet, the receive logic 32 indicates with the word whether this packet is a candidate for fast-path processing. FIG. 4A depicts the case in which
 15 the packet is not a fast-path candidate, in which case the CPD 30 sends the validated headers and data from memory 60 to data link layer 36 along an internal bus for processing by the host CPU, as shown by arrow 56. The packet is processed by the host protocol stack 44 of data link 36, network 38, transport 40 and session 42 layers, and data (D) 63 from the packet may then be sent to storage 35, as shown by arrow 65.

FIG. 4B, depicts the case in which the receive logic 32 of the CPD determines that a message packet is a candidate for fast-path processing, for example by deriving from the packet's headers that the packet belongs to a TCP/IP, TTCP/IP or SPX/IPX message. A processor 55 in the CPD 30 then checks to see whether the word that summarizes the fast-path candidate matches a CCB held in a cache 62. Upon finding no match for this packet, the
 25 CPD sends the validated packet from memory 60 to the host protocol stack 44 for processing. Host stack 44 may use this packet to create a connection context for the message, including finding and reserving a destination for data from the message associated with the packet, the context taking the form of a CCB. The present embodiment employs a single specialized host stack 44 for processing both fast-path and non-fast-path candidates, while in an
 30 embodiment described below fast-path candidates are processed by a different host stack than non-fast-path candidates. Some data (D1) 66 from that initial packet may optionally be sent to the destination in storage 35, as shown by arrow 68. The CCB is then sent to the CPD 30 to be saved in cache 62, as shown by arrow 64. For a traditional connection-based message

WO 00/13091

PCT/US98/24943

such as typified by TCP/IP, the initial packet may be part of a connection initialization dialogue that transpires between hosts before the CCB is created and passed to the CPD 30.

Referring now to FIG. 4C, when a subsequent packet from the same connection as the initial packet is received from the network 25 by CPD 30, the packet headers and data are
 5 validated by the receive logic 32, and the headers are parsed to create a summary of the message packet and a hash for finding a corresponding CCB, the summary and hash contained in a word or words. The word or words are temporarily stored in memory 60 along with the packet. The processor 55 checks for a match between the hash and each CCB that is stored in the cache 62 and, finding a match, sends the data (D2) 70 via a fast-path directly to
 10 the destination in storage 35, as shown by arrow 72, bypassing the session layer 42, transport layer 40, network layer 38 and data link layer 36. The remaining data packets from the message can also be sent by DMA directly to storage, avoiding the relatively slow protocol layer processing and repeated copying by the CPU stack 44.

FIG. 4D shows the procedure for handling the rare instance when a message for
 15 which a fast-path connection has been established, such as shown in FIG. 4C, has a packet that is not easily handled by the CPD. In this case the packet is sent to be processed by the protocol stack 44, which is handed the CCB for that message from cache 62 via a control dialogue with the CPD, as shown by arrow 76, signaling to the CPU to take over processing of that message. Slow-path processing by the protocol stack then results in data (D3) 80 from
 20 the packet being sent, as shown by arrow 82, to storage 35. Once the packet has been processed and the error situation corrected, the CCB can be handed back via a control dialogue to the cache 62, so that payload data from subsequent packets of that message can again be sent via the fast-path of the CPD 30. Thus the CPU and CPD together decide whether a given message is to be processed according to fast-path hardware processing or
 25 more conventional software processing by the CPU.

Transmission of a message from the host 20 to the network 25 for delivery to remote host 22 also can be processed by either sequential protocol software processing via the CPU or accelerated hardware processing via the CPD 30, as shown in FIG. 5. A message (M) 90 that is selected by CPU 28 from storage 35 can be sent to session layer 42 for processing by
 30 stack 44, as shown by arrows 92 and 96. For the situation in which a connection exists and the CPD 30 already has an appropriate CCB for the message, however, data packets can bypass host stack 44 and be sent by DMA directly to memory 60, with the processor 55 adding to each data packet a single header containing all the appropriate protocol layers, and sending the resulting packets to the network 25 for transmission to remote host 22. This fast-

WO 00/13091

PCT/US98/24943

path transmission can greatly accelerate processing for even a single packet, with the acceleration multiplied for a larger message.

A message for which a fast-path connection is not extant thus may benefit from creation of a CCB with appropriate control and state information for guiding fast-path transmission. For a traditional connection-based message, such as typified by TCP/IP or SPX/IPX, the CCB is created during connection initialization dialogue. For a quick-connection message, such as typified by TTCP/IP, the CCB can be created with the same transaction that transmits payload data. In this case, the transmission of payload data may be a reply to a request that was used to set up the fast-path connection. In any case, the CCB provides protocol and status information regarding each of the protocol layers, including which user is involved and storage space for per-transfer information. The CCB is created by protocol stack 44, which then passes the CCB to the CPD 30 by writing to a command register of the CPD, as shown by arrow 98. Guided by the CCB, the processor 55 moves network frame-sized portions of the data from the source in host memory 35 into its own memory 60 using DMA, as depicted by arrow 99. The processor 55 then prepends appropriate headers and checksums to the data portions, and transmits the resulting frames to the network 25, consistent with the restrictions of the associated protocols. After the CPD 30 has received an acknowledgement that all the data has reached its destination, the CPD will then notify the host 35 by writing to a response buffer.

Thus, fast-path transmission of data communications also relieves the host CPU of per-frame processing. A vast majority of data transmissions can be sent to the network by the fast-path. Both the input and output fast-paths attain a huge reduction in interrupts by functioning at an upper layer level, i.e., session level or higher, and interactions between the network microprocessor and the host occur using the full transfer sizes which that upper layer wishes to make. For fast-path communications, an interrupt only occurs (at the most) at the beginning and end of an entire upper-layer message transaction, and there are no interrupts for the sending or receiving of each lower layer portion or packet of that transaction.

A simplified intelligent network interface card (INIC) 150 is shown in FIG. 6 to provide a network interface for a host 152. Hardware logic 171 of the INIC 150 is connected to a network 155, with a peripheral bus (PCI) 157 connecting the INIC and host. The host 152 in this embodiment has a TCP/IP protocol stack, which provides a slow-path 158 for sequential software processing of message frames received from the network 155. The host 152 protocol stack includes a data link layer 160, network layer 162, a transport layer 164 and an application layer 166, which provides a source or destination 168 for the communication

WO 00/13091

PCT/US98/24943

data in the host 152. Other layers which are not shown, such as session and presentation layers, may also be included in the host stack 152, and the source or destination may vary depending upon the nature of the data and may actually be the application layer.

The INIC 150 has a network processor 170 which chooses between processing
 5 messages along a slow-path 158 that includes the protocol stack of the host, or along a fast-path 159 that bypasses the protocol stack of the host. Each received packet is processed on the fly by hardware logic 171 contained in INIC 150, so that all of the protocol headers for a packet can be processed without copying, moving or storing the data between protocol layers. The hardware logic 171 processes the headers of a given packet at one time as packet bytes
 10 pass through the hardware, by categorizing selected header bytes. Results of processing the selected bytes help to determine which other bytes of the packet are categorized, until a summary of the packet has been created, including checksum validations. The processed headers and data from the received packet are then stored in INIC storage 185, as well as the word or words summarizing the headers and status of the packet. For a network storage
 15 configuration, the INIC 150 may be connected to a peripheral storage device such as a disk drive which has an IDE, SCSI or similar interface, with a file cache for the storage device residing on the memory 185 of the INIC 150. Several such network interfaces may exist for a host, with each interface having an associated storage device.

The hardware processing of message packets received by INIC 150 from network 155
 20 is shown in more detail in FIG. 7. A received message packet first enters a media access controller 172, which controls INIC access to the network and receipt of packets and can provide statistical information for network protocol management. From there, data flows one byte at a time into an assembly register 174, which in this example is 128 bits wide. The data is categorized by a fly-by sequencer 178, as will be explained in more detail with regard to
 25 FIG. 8, which examines the bytes of a packet as they fly by, and generates status from those bytes that will be used to summarize the packet. The status thus created is merged with the data by a multiplexor 180 and the resulting data stored in SRAM 182. A packet control sequencer 176 oversees the fly-by sequencer 178, examines information from the media access controller 172, counts the bytes of data, generates addresses, moves status and
 30 manages the movement of data from the assembly register 174 to SRAM 182 and eventually DRAM 188. The packet control sequencer 176 manages a buffer in SRAM 182 via SRAM controller 183, and also indicates to a DRAM controller 186 when data needs to be moved from SRAM 182 to a buffer in DRAM 188. Once data movement for the packet has been completed and all the data has been moved to the buffer in DRAM 188, the packet control

WO 00/13091

PCT/US98/24943

sequencer 176 will move the status that has been generated in the fly-by sequencer 178 out to the SRAM 182 and to the beginning of the DRAM 188 buffer to be prepended to the packet data. The packet control sequencer 176 then requests a queue manager 184 to enter a receive buffer descriptor into a receive queue, which in turn notifies the processor 170 that the packet
5 has been processed by hardware logic 171 and its status summarized.

FIG. 8 shows that the fly-by sequencer 178 has several tiers, with each tier generally focusing on a particular portion of the packet header and thus on a particular protocol layer, for generating status pertaining to that layer. The fly-by sequencer 178 in this embodiment includes a media access control sequencer 191, a network sequencer 192, a transport
10 sequencer 194 and a session sequencer 195. Sequencers pertaining to higher protocol layers can additionally be provided. The fly-by sequencer 178 is reset by the packet control sequencer 176 and given pointers by the packet control sequencer that tell the fly-by sequencer whether a given byte is available from the assembly register 174. The media access control sequencer 191 determines, by looking at bytes 0-5, that a packet is addressed
15 to host 152 rather than or in addition to another host. Offsets 12 and 13 of the packet are also processed by the media access control sequencer 191 to determine the type field, for example whether the packet is Ethernet or 802.3. If the type field is Ethernet those bytes also tell the media access control sequencer 191 the packet's network protocol type. For the 802.3 case, those bytes instead indicate the length of the entire frame, and the media access control
20 sequencer 191 will check eight bytes further into the packet to determine the network layer type.

For most packets the network sequencer 192 validates that the header length received has the correct length, and checksums the network layer header. For fast-path candidates the network layer header is known to be IP or IPX from analysis done by the media access
25 control sequencer 191. Assuming for example that the type field is 802.3 and the network protocol is IP, the network sequencer 192 analyzes the first bytes of the network layer header, which will begin at byte 22, in order to determine IP type. The first bytes of the IP header will be processed by the network sequencer 192 to determine what IP type the packet involves. Determining that the packet involves, for example, IP version 4, directs further
30 processing by the network sequencer 192, which also looks at the protocol type located ten bytes into the IP header for an indication of the transport header protocol of the packet. For example, for IP over Ethernet, the IP header begins at offset 14, and the protocol type byte is offset 23, which will be processed by network logic to determine whether the transport layer protocol is TCP, for example. From the length of the network layer header, which is

WO 00/13091

PCT/US98/24943

typically 20-40 bytes, network sequencer 192 determines the beginning of the packet's transport layer header for validating the transport layer header. Transport sequencer 194 may generate checksums for the transport layer header and data, which may include information from the IP header in the case of TCP at least.

5 Continuing with the example of a TCP packet, transport sequencer 194 also analyzes the first few bytes in the transport layer portion of the header to determine, in part, the TCP source and destination ports for the message, such as whether the packet is NetBios or other protocols. Byte 12 of the TCP header is processed by the transport sequencer 194 to determine and validate the TCP header length. Byte 13 of the TCP header contains flags that
10 may, aside from ack flags and push flags, indicate unexpected options, such as reset and fin, that may cause the processor to categorize this packet as an exception. TCP offset bytes 16 and 17 are the checksum, which is pulled out and stored by the hardware logic 171 while the rest of the frame is validated against the checksum.

Session sequencer 195 determines the length of the session layer header, which in the
15 case of NetBios is only four bytes, two of which tell the length of the NetBios payload data, but which can be much larger for other protocols. The session sequencer 195 can also be used to categorize the type of message as read or write, for example, for which the fast-path may be particularly beneficial. Further upper layer logic processing, depending upon the message type, can be performed by the hardware logic 171 of packet control sequencer 176
20 and fly-by sequencer 178. Thus hardware logic 171 intelligently directs hardware processing of the headers by categorization of selected bytes from a single stream of bytes, with the status of the packet being built from classifications determined on the fly. Once the packet control sequencer 176 detects that all of the packet has been processed by the fly-by sequencer 178, the packet control sequencer 176 adds the status information generated by the
25 fly-by sequencer 178 and any status information generated by the packet control sequencer 176, and prepends (adds to the front) that status information to the packet, for convenience in handling the packet by the processor 170. The additional status information generated by the packet control sequencer 176 includes media access controller 172 status information and any errors discovered, or data overflow in either the assembly register or DRAM buffer, or other
30 miscellaneous information regarding the packet. The packet control sequencer 176 also stores entries into a receive buffer queue and a receive statistics queue via the queue manager 184.

An advantage of processing a packet by hardware logic 171 is that the packet does not, in contrast with conventional sequential software protocol processing, have to be stored,

WO 00/13091

PCT/US98/24943

5 moved, copied or pulled from storage for processing each protocol layer header, offering dramatic increases in processing efficiency and savings in processing time for each packet. The packets can be processed at the rate bits are received from the network, for example 100 megabits/second for a 100 baseT connection. The time for categorizing a packet received at this rate and having a length of sixty bytes is thus about 5 microseconds. The total time for processing this packet with the hardware logic 171 and sending packet data to its host destination via the fast-path may be about 16 microseconds or less, assuming a 66 MHz PCI bus, whereas conventional software protocol processing by a 300 MHz Pentium II® processor may take as much as 200 microseconds in a busy device. More than an order of magnitude
10 decrease in processing time can thus be achieved with fast-path 159 in comparison with a high-speed CPU employing conventional sequential software protocol processing, demonstrating the dramatic acceleration provided by processing the protocol headers by the hardware logic 171 and processor 170, without even considering the additional time savings afforded by the reduction in CPU interrupts and host bus bandwidth savings.

15 The processor 170 chooses, for each received message packet held in storage 185, whether that packet is a candidate for the fast-path 159 and, if so, checks to see whether a fast-path has already been set up for the connection that the packet belongs to. To do this, the processor 170 first checks the header status summary to determine whether the packet headers are of a protocol defined for fast-path candidates. If not, the processor 170
20 commands DMA controllers in the INIC 150 to send the packet to the host for slow-path 158 processing. Even for a slow-path 158 processing of a message, the INIC 150 thus performs initial procedures such as validation and determination of message type, and passes the validated message at least to the data link layer 160 of the host.

For fast-path 159 candidates, the processor 170 checks to see whether the header
25 status summary matches a CCB held by the INIC. If so, the data from the packet is sent along fast-path 159 to the destination 168 in the host. If the fast-path 159 candidate's packet summary does not match a CCB held by the INIC, the packet may be sent to the host 152 for slow-path processing to create a CCB for the message. Employment of the fast-path 159 may also not be needed or desirable for the case of fragmented messages or other complexities.
30 For the vast majority of messages, however, the INIC fast-path 159 can greatly accelerate message processing. The INIC 150 thus provides a single state machine processor 170 that decides whether to send data directly to its destination, based upon information gleaned on the fly, as opposed to the conventional employment of a state machine in each of several protocol layers for determining the destiny of a given packet.

WO 00/13091

PCT/US98/24943

In processing an indication or packet received at the host 152, a protocol driver of the host selects the processing route based upon whether the indication is fast-path or slow-path. A TCP/IP or SPX/IPX message has a connection that is set up from which a CCB is formed by the driver and passed to the INIC for matching with and guiding the fast-path packet to the connection destination 168. For a TTCP/IP message, the driver can create a connection context for the transaction from processing an initial request packet, including locating the message destination 168, and then passing that context to the INIC in the form of a CCB for providing a fast-path for a reply from that destination. A CCB includes connection and state information regarding the protocol layers and packets of the message. Thus a CCB can include source and destination media access control (MAC) addresses, source and destination IP or IPX addresses, source and destination TCP or SPX ports, TCP variables such as timers, receive and transmit windows for sliding window protocols, and information indicating the session layer protocol.

Caching the CCBs in a hash table in the INIC provides quick comparisons with words summarizing incoming packets to determine whether the packets can be processed via the fast-path 159, while the full CCBs are also held in the INIC for processing. Other ways to accelerate this comparison include software processes such as a B-tree or hardware assists such as a content addressable memory (CAM). When INIC microcode or comparator circuits detect a match with the CCB, a DMA controller places the data from the packet in the destination 168, without any interrupt by the CPU, protocol processing or copying. Depending upon the type of message received, the destination of the data may be the session, presentation or application layers, or a file buffer cache in the host 152.

FIG. 9 shows an INIC 200 connected to a host 202 that is employed as a file server. This INIC provides a network interface for several network connections employing the 802.3u standard, commonly known as Fast Ethernet. The INIC 200 is connected by a PCI bus 205 to the server 202, which maintains a TCP/IP or SPX/IPX protocol stack including MAC layer 212, network layer 215, transport layer 217 and application layer 220, with a source/destination 222 shown above the application layer, although as mentioned earlier the application layer can be the source or destination. The INIC is also connected to network lines 210, 240, 242 and 244, which are preferably Fast Ethernet, twisted pair, fiber optic, coaxial cable or other lines each allowing data transmission of 100 Mb/s, while faster and slower data rates are also possible. Network lines 210, 240, 242 and 244 are each connected to a dedicated row of hardware circuits which can each validate and summarize message packets received from their respective network line. Thus line 210 is connected with a first

WO 00/13091

PCT/US98/24943

horizontal row of sequencers 250, line 240 is connected with a second horizontal row of sequencers 260, line 242 is connected with a third horizontal row of sequencers 262 and line 244 is connected with a fourth horizontal row of sequencers 264. After a packet has been validated and summarized by one of the horizontal hardware rows it is stored along with its status summary in storage 270.

A network processor 230 determines, based on that summary and a comparison with any CCBs stored in the INIC 200, whether to send a packet along a slow-path 231 for processing by the host. A large majority of packets can avoid such sequential processing and have their data portions sent by DMA along a fast-path 237 directly to the data destination 222 in the server according to a matching CCB. Similarly, the fast-path 237 provides an avenue to send data directly from the source 222 to any of the network lines by processor 230 division of the data into packets and addition of full headers for network transmission, again minimizing CPU processing and interrupts. For clarity only horizontal sequencer 250 is shown active; in actuality each of the sequencer rows 250, 260, 262 and 264 offers full duplex communication, concurrently with all other sequencer rows. The specialized INIC 200 is much faster at working with message packets than even advanced general-purpose host CPUs that processes those headers sequentially according to the software protocol stack.

One of the most commonly used network protocols for large messages such as file transfers is server message block (SMB) over TCP/IP. SMB can operate in conjunction with redirector software that determines whether a required resource for a particular operation, such as a printer or a disk upon which a file is to be written, resides in or is associated with the host from which the operation was generated or is located at another host connected to the network, such as a file server. SMB and server/redirector are conventionally serviced by the transport layer; in the present invention SMB and redirector can instead be serviced by the INIC. In this case, sending data by the DMA controllers from the INIC buffers when receiving a large SMB transaction may greatly reduce interrupts that the host must handle. Moreover, this DMA generally moves the data to its final destination in the file device cache. An SMB transmission of the present invention follows essentially the reverse of the above described SMB receive, with data transferred from the host to the INIC and stored in buffers, while the associated protocol headers are prepended to the data in the INIC, for transmission via a network line to a remote host. Processing by the INIC of the multiple packets and multiple TCP, IP, NetBios and SMB protocol layers via custom hardware and without repeated interrupts of the host can greatly increase the speed of transmitting an SMB message to a network line.

WO 00/13091

PCT/US98/24943

As shown in FIG. 10, for controlling whether a given message is processed by the host 202 or by the INIC 200, a message command driver 300 may be installed in host 202 to work in concert with a host protocol stack 310. The command driver 300 can intervene in message reception or transmittal, create CCBs and send or receive CCBs from the INIC 200, so that functioning of the INIC, aside from improved performance, is transparent to a user. Also shown is an INIC memory 304 and an INIC miniport driver 306, which can direct message packets received from network 210 to either the conventional protocol stack 310 or the command protocol stack 300, depending upon whether a packet has been labeled as a fast-path candidate. The conventional protocol stack 310 has a data link layer 312, a network layer 314 and a transport layer 316 for conventional, lower layer processing of messages that are not labeled as fast-path candidates and therefore not processed by the command stack 300. Residing above the lower layer stack 310 is an upper layer 318, which represents a session, presentation and/or application layer, depending upon the message communicated. The command driver 300 similarly has a data link layer 320, a network layer 322 and a transport layer 325.

The driver 300 includes an upper layer interface 330 that determines, for transmission of messages to the network 210, whether a message transmitted from the upper layer 318 is to be processed by the command stack 300 and subsequently the INIC fast-path, or by the conventional stack 310. When the upper layer interface 330 receives an appropriate message from the upper layer 318 that would conventionally be intended for transmission to the network after protocol processing by the protocol stack of the host, the message is passed to driver 300. The INIC then acquires network-sized portions of the message data for that transmission via INIC DMA units, prepends headers to the data portions and sends the resulting message packets down the wire. Conversely, in receiving a TCP, TTCP, SPX or similar message packet from the network 210 to be used in setting up a fast-path connection, miniport driver 306 diverts that message packet to command driver 300 for processing. The driver 300 processes the message packet to create a context for that message, with the driver 302 passing the context and command instructions back to the INIC 200 as a CCB for sending data of subsequent messages for the same connection along a fast-path. Hundreds of TCP, TTCP, SPX or similar CCB connections may be held indefinitely by the INIC, although a least recently used (LRU) algorithm is employed for the case when the INIC cache is full. The driver 300 can also create a connection context for a TTCP request which is passed to the INIC 200 as a CCB, allowing fast-path transmission of a TTCP reply to the request. A

WO 00/13091

PCT/US98/24943

message having a protocol that is not accelerated can be processed conventionally by protocol stack 310.

FIG. 11 shows a TCP/IP implementation of command driver software for Microsoft® protocol messages. A conventional host protocol stack 350 includes MAC layer 353, IP layer 355 and TCP layer 358. A command driver 360 works in concert with the host stack 350 to process network messages. The command driver 360 includes a MAC layer 363, an IP layer 366 and an Alacritech TCP (ATCP) layer 373. The conventional stack 350 and command driver 360 share a network driver interface specification (NDIS) layer 375, which interacts with the INIC miniport driver 306. The INIC miniport driver 306 sorts receive indications for processing by either the conventional host stack 350 or the ATCP driver 360. A TDI filter driver and upper layer interface 380 similarly determines whether messages sent from a TDI user 382 to the network are diverted to the command driver and perhaps to the fast-path of the INIC, or processed by the host stack.

FIG. 12 depicts a typical SMB exchange between a client 190 and server 290, both of which have communication devices of the present invention, the communication devices each holding a CCB defining their connection for fast-path movement of data. The client 190 includes INIC 150, 802.3 compliant data link layer 160, IP layer 162, TCP layer 164, NetBios layer 166, and SMB layer 168. The client has a slow-path 157 and fast-path 159 for communication processing. Similarly, the server 290 includes INIC 200, 802.3 compliant data link layer 212, IP layer 215, TCP layer 217, NetBios layer 220, and SMB 222. The server is connected to network lines 240, 242 and 244, as well as line 210 which is connected to client 190. The server also has a slow-path 231 and fast-path 237 for communication processing.

Assuming that the client 190 wishes to read a 100KB file on the server 290, the client may begin by sending a Read Block Raw (RBR) SMB command across network 210 requesting the first 64 KB of that file on the server 290. The RBR command may be only 76 bytes, for example, so the INIC 200 on the server will recognize the message type (SMB) and relatively small message size, and send the 76 bytes directly via the fast-path to NetBios of the server. NetBios will give the data to SMB, which processes the Read request and fetches the 64KB of data into server data buffers. SMB then calls NetBios to send the data, and NetBios outputs the data for the client. In a conventional host, NetBios would call TCP output and pass 64 KB to TCP, which would divide the data into 1460 byte segments and output each segment via IP and eventually MAC (slow-path 231). In the present case, the 64KB data goes to the ATCP driver along with an indication regarding the client-server SMB

WO 00/13091

PCT/US98/24943

connection, which indicates a CCB held by the INIC. The INIC 200 then proceeds to DMA 1460 byte segments from the host buffers, add the appropriate headers for TCP, IP and MAC at one time, and send the completed packets on the network 210 (fast-path 237). The INIC 200 will repeat this until the whole 64KB transfer has been sent. Usually after receiving
 5 acknowledgement from the client that the 64KB has been received, the INIC will then send the remaining 36KB also by the fast-path 237.

With INIC 150 operating on the client 190 when this reply arrives, the INIC 150 recognizes from the first frame received that this connection is receiving fast-path 159 processing (TCP/IP, NetBios, matching a CCB), and the ATCP may use this first frame to
 10 acquire buffer space for the message. This latter case is done by passing the first 128 bytes of the NetBios portion of the frame via the ATCP fast-path directly to the host NetBios; that will give NetBios/SMB all of the frame's headers. NetBios/SMB will analyze these headers, realize by matching with a request ID that this is a reply to the original RawRead connection, and give the ATCP a 64K list of buffers into which to place the data. At this stage only one
 15 frame has arrived, although more may arrive while this processing is occurring. As soon as the client buffer list is given to the ATCP, it passes that transfer information to the INIC 150, and the INIC 150 starts DMAing any frame data that has accumulated into those buffers.

FIG. 13 provides a simplified diagram of the INIC 200, which combines the functions of a network interface controller and a protocol processor in a single ASIC chip 400. The
 20 INIC 200 in this embodiment offers a full-duplex, four channel, 10/100-Megabit per second (Mbps) intelligent network interface controller that is designed for high speed protocol processing for server applications. Although designed specifically for server applications, the INIC 200 can be connected to personal computers, workstations, routers or other hosts anywhere that TCP/IP, TTCP/IP or SPX/IPX protocols are being utilized.

25 The INIC 200 is connected with four network lines 210, 240, 242 and 244, which may transport data along a number of different conduits, such as twisted pair, coaxial cable or optical fiber, each of the connections providing a media independent interface (MII) via commercially available physical layer chips, such as model 80220/80221 Ethernet Media Interface Adapter from SEEQ Technology Incorporated, 47200 Bayside Parkway, Fremont,
 30 CA 94538. The lines preferably are 802.3 compliant and in connection with the INIC constitute four complete Ethernet nodes, the INIC supporting 10Base-T, 10Base-T2, 100Base-TX, 100Base-FX and 100Base-T4 as well as future interface standards. Physical layer identification and initialization is accomplished through host driver initialization routines. The connection between the network lines 210, 240, 242 and 244 and the INIC 200

WO 00/13091

PCT/US98/24943

is controlled by MAC units MAC-A 402, MAC-B 404, MAC-C 406 and MAC-D 408 which contain logic circuits for performing the basic functions of the MAC sublayer, essentially controlling when the INIC accesses the network lines 210, 240, 242 and 244. The MAC units 402-408 may act in promiscuous, multicast or unicast modes, allowing the INIC to function
5 as a network monitor, receive broadcast and multicast packets and implement multiple MAC addresses for each node. The MAC units 402-408 also provide statistical information that can be used for simple network management protocol (SNMP).

The MAC units 402, 404, 406 and 408 are each connected to a transmit and receive sequencer, XMT & RCV-A 418, XMT & RCV-B 420, XMT & RCV-C 422 and XMT &
10 RCV-D 424, by wires 410, 412, 414 and 416, respectively. Each of the transmit and receive sequencers can perform several protocol processing steps on the fly as message frames pass through that sequencer. In combination with the MAC units, the transmit and receive sequencers 418-422 can compile the packet status for the data link, network, transport, session and, if appropriate, presentation and application layer protocols in hardware, greatly
15 reducing the time for such protocol processing compared to conventional sequential software engines. The transmit and receive sequencers 410-414 are connected, by lines 426, 428, 430 and 432 to an SRAM and DMA controller 444, which includes DMA controllers 438 and SRAM controller 442. Static random access memory (SRAM) buffers 440 are coupled with SRAM controller 442 by line 441. The SRAM and DMA controllers 444 interact across line
20 446 with external memory control 450 to send and receive frames via external memory bus 455 to and from dynamic random access memory (DRAM) buffers 460, which is located adjacent to the IC chip 400. The DRAM buffers 460 may be configured as 4 MB, 8 MB, 16 MB or 32 MB, and may optionally be disposed on the chip. The SRAM and DMA controllers 444 are connected via line 464 to a PCI Bus Interface Unit (BIU) 468, which
25 manages the interface between the INIC 200 and the PCI interface bus 257. The 64-bit, multiplexed BIU 468 provides a direct interface to the PCI bus 257 for both slave and master functions. The INIC 200 is capable of operating in either a 64-bit or 32-bit PCI environment, while supporting 64-bit addressing in either configuration.

A microprocessor 470 is connected by line 472 to the SRAM and DMA controllers
30 444, and connected via line 475 to the PCI BIU 468. Microprocessor 470 instructions and register files reside in an on chip control store 480, which includes a writable on-chip control store (WCS) of SRAM and a read only memory (ROM), and is connected to the microprocessor by line 477. The microprocessor 470 offers a programmable state machine which is capable of processing incoming frames, processing host commands, directing

WO 00/13091

PCT/US98/24943

network traffic and directing PCI bus traffic. Three processors are implemented using shared hardware in a three level pipelined architecture that launches and completes a single instruction for every clock cycle. A receive processor 482 is primarily used for receiving communications while a transmit processor 484 is primarily used for transmitting
5 communications in order to facilitate full duplex communication, while a utility processor 486 offers various functions including overseeing and controlling PCI register access.

The instructions for the three processors 482, 484 and 486 reside in the on-chip control-store 480. Thus the functions of the three processors can be easily redefined, so that the microprocessor 470 can adapted for a given environment. For instance, the amount of
10 processing required for receive functions may outweigh that required for either transmit or utility functions. In this situation, some receive functions may be performed by the transmit processor 484 and/or the utility processor 486. Alternatively, an additional level of pipelining can be created to yield four or more virtual processors instead of three, with the additional level devoted to receive functions.

15 The INIC 200 in this embodiment can support up to 256 CCBs which are maintained in a table in the DRAM 460. There is also, however, a CCB index in hash order in the SRAM 440 to save sequential searching. Once a hash has been generated, the CCB is cached in SRAM, with up to sixteen cached CCBs in SRAM in this example. Allocation of the sixteen CCBs cached in SRAM is handled by a least recently used register, described below.
20 These cache locations are shared between the transmit 484 and receive 486 processors so that the processor with the heavier load is able to use more cache buffers. There are also eight header buffers and eight command buffers to be shared between the sequencers. A given header or command buffer is not statically linked to a specific CCB buffer, as the link is dynamic on a per-frame basis.

25 FIG. 14 shows an overview of the pipelined microprocessor 470, in which instructions for the receive, transmit and utility processors are executed in three alternating phases according to Clock increments I, II and III, the phases corresponding to each of the pipeline stages. Each phase is responsible for different functions, and each of the three processors occupies a different phase during each Clock increment. Each processor usually operates
30 upon a different instruction stream from the control store 480, and each carries its own program counter and status through each of the phases.

In general, a first instruction phase 500 of the pipelined microprocessors completes an instruction and stores the result in a destination operand, fetches the next instruction, and stores that next instruction in an instruction register. A first register set 490 provides a

WO 00/13091

PCT/US98/24943

number of registers including the instruction register, and a set of controls 492 for first register set provides the controls for storage to the first register set 490. Some items pass through the first phase without modification by the controls 492, and instead are simply copied into the first register set 490 or a RAM file register 533. A second instruction phase

5 560 has an instruction decoder and operand multiplexer 498 that generally decodes the instruction that was stored in the instruction register of the first register set 490 and gathers any operands which have been generated, which are then stored in a decode register of a second register set 496. The first register set 490, second register set 496 and a third register set 501, which is employed in a third instruction phase 600, include many of the same

10 registers, as will be seen in the more detailed views of FIGs. 15A-C. The instruction decoder and operand multiplexer 498 can read from two address and data ports of the RAM file register 533, which operates in both the first phase 500 and second phase 560. A third phase 600 of the processor 470 has an arithmetic logic unit (ALU) 602 which generally performs any ALU operations on the operands from the second register set, storing the results in a

15 results register included in the third register set 501. A stack exchange 608 can reorder register stacks, and a queue manager 503 can arrange queues for the processor 470, the results of which are stored in the third register set.

The instructions continue with the first phase then following the third phase, as depicted by a circular pipeline 505. Note that various functions have been distributed across

20 the three phases of the instruction execution in order to minimize the combinatorial delays within any given phase. With a frequency in this embodiment of 66 MHz, each Clock increment takes 15 nanoseconds to complete, for a total of 45 nanoseconds to complete one instruction for each of the three processors. The rotating instruction phases are depicted in more detail in FIGs. 15A-C, in which each phase is shown in a different figure.

25 More particularly, FIG. 15A shows some specific hardware functions of the first phase 500, which generally includes the first register set 490 and related controls 492. The controls for the first register set 492 includes an SRAM control 502, which is a logical control for loading address and write data into SRAM address and data registers 520. Thus the output of the ALU 602 from the third phase 600 may be placed by SRAM control 502

30 into an address register or data register of SRAM address and data registers 520. A load control 504 similarly provides controls for writing a context for a file to file context register 522, and another load control 506 provides controls for storing a variety of miscellaneous data to flip-flop registers 525. ALU condition codes, such as whether a carried bit is set, get clocked into ALU condition codes register 528 without an operation performed in the first

WO 00/13091

PCT/US98/24943

phase 500. Flag decodes 508 can perform various functions, such as setting locks, that get stored in flag registers 530.

The RAM file register 533 has a single write port for addresses and data and two read ports for addresses and data, so that more than one register can be read from at one time. As noted above, the RAM file register 533 essentially straddles the first and second phases, as it is written in the first phase 500 and read from in the second phase 560. A control store instruction 510 allows the reprogramming of the processors due to new data in from the control store 480, not shown in this figure, the instructions stored in an instruction register 535. The address for this is generated in a fetch control register 511, which determines which address to fetch, the address stored in fetch address register 538. Load control 515 provides instructions for a program counter 540, which operates much like the fetch address for the control store. A last-in first-out stack 544 of three registers is copied to the first register set without undergoing other operations in this phase. Finally, a load control 517 for a debug address 548 is optionally included, which allows correction of errors that may occur.

FIG. 15B depicts the second microprocessor phase 560, which includes reading addresses and data out of the RAM file register 533. A scratch SRAM 565 is written from SRAM address and data register 520 of the first register set, which includes a register that passes through the first two phases to be incremented in the third. The scratch SRAM 565 is read by the instruction decoder and operand multiplexer 498, as are most of the registers from the first register set, with the exception of the stack 544, debug address 548 and SRAM address and data register mentioned above. The instruction decoder and operand multiplexer 498 looks at the various registers of set 490 and SRAM 565, decodes the instructions and gathers the operands for operation in the next phase, in particular determining the operands to provide to the ALU 602 below. The outcome of the instruction decoder and operand multiplexer 498 is stored to a number of registers in the second register set 496, including ALU operands 579 and 582, ALU condition code register 580, and a queue channel and command 587 register, which in this embodiment can control thirty-two queues. Several of the registers in set 496 are loaded fairly directly from the instruction register 535 above without substantial decoding by the decoder 498, including a program control 590, a literal field 589, a test select 584 and a flag select 585. Other registers such as the file context 522 of the first phase 500 are always stored in a file context 577 of the second phase 560, but may also be treated as an operand that is gathered by the multiplexer 572. The stack registers 544 are simply copied in stack register 594. The program counter 540 is incremented 568 in this phase and stored in register 592. Also incremented 570 is the optional debug address 548,

WO 00/13091

PCT/US98/24943

and a load control 575 may be fed from the pipeline 505 at this point in order to allow error control in each phase, the result stored in debug address 598.

FIG. 15C depicts the third microprocessor phase 600, which includes ALU and queue operations. The ALU 602 includes an adder, priority encoders and other standard logic functions. Results of the ALU are stored in registers ALU output 618, ALU condition codes 620 and destination operand results 622. A file context register 616, flag select register 626 and literal field register 630 are simply copied from the previous phase 560. A test multiplexer 604 is provided to determine whether a conditional jump results in a jump, with the results stored in a test results register 624. The test multiplexer 604 may instead be performed in the first phase 500 along with similar decisions such as fetch control 511. A stack exchange 608 shifts a stack up or down by fetching a program counter from stack 594 or putting a program counter onto that stack, results of which are stored in program control 634, program counter 638 and stack 640 registers. The SRAM address may optionally be incremented in this phase 600. Another load control 610 for another debug address 642 may be forced from the pipeline 505 at this point in order to allow error control in this phase also. A QRAM & QALU 606, shown together in this figure, read from the queue channel and command register 587, store in SRAM and rearrange queues, adding or removing data and pointers as needed to manage the queues of data, sending results to the test multiplexer 604 and a queue flags and queue address register 628. Thus the QRAM & QALU 606 assume the duties of managing queues for the three processors, a task conventionally performed sequentially by software on a CPU, the queue manager 606 instead providing accelerated and substantially parallel hardware queuing.

FIG. 16 depicts two of the thirty-two hardware queues that are managed by the queue manager 606, with each of the queues having an SRAM head, an SRAM tail and the ability to queue information in a DRAM body as well, allowing expansion and individual configuration of each queue. Thus FIFO 700 has SRAM storage units, 705, 707, 709 and 711, each containing eight bytes for a total of thirty-two bytes, although the number and capacity of these units may vary in other embodiments. Similarly, FIFO 702 has SRAM storage units 713, 715, 717 and 719. SRAM units 705 and 707 are the head of FIFO 700 and units 709 and 711 are the tail of that FIFO, while units 713 and 715 are the head of FIFO 702 and units 717 and 719 are the tail of that FIFO. Information for FIFO 700 may be written into head units 705 or 707, as shown by arrow 722, and read from tail units 711 or 709, as shown by arrow 725. A particular entry, however, may be both written to and read from head units 705 or 707, or may be both written to and read from tail units 709 or 711, minimizing data

WO 00/13091

PCT/US98/24943

movement and latency. Similarly, information for FIFO 702 is typically written into head units 713 or 715, as shown by arrow 733, and read from tail units 717 or 719, as shown by arrow 739, but may instead be read from the same head or tail unit to which it was written.

5 The SRAM FIFOS 700 and 702 are both connected to DRAM 460, which allows virtually unlimited expansion of those FIFOS to handle situations in which the SRAM head and tail are full. For example a first of the thirty-two queues, labeled Q-zero, may queue an entry in DRAM 460, as shown by arrow 727, by DMA units acting under direction of the queue manager, instead of being queued in the head or tail of FIFO 700. Entries stored in
10 DRAM 460 return to SRAM unit 709, as shown by arrow 730, extending the length and fall-through time of that FIFO. Diversion from SRAM to DRAM is typically reserved for when the SRAM is full, since DRAM is slower and DMA movement causes additional latency. Thus Q-zero may comprise the entries stored by queue manager 606 in both the FIFO 700 and the DRAM 460. Likewise, information bound for FIFO 702, which may correspond to
15 Q-twenty-seven, for example, can be moved by DMA into DRAM 460, as shown by arrow 735. The capacity for queuing in cost-effective albeit slower DRAM 460 is user-definable during initialization, allowing the queues to change in size as desired. Information queued in DRAM 460 is returned to SRAM unit 717, as shown by arrow 737.

 Status for each of the thirty-two hardware queues is conveniently maintained in and
20 accessed from a set 740 of four, thirty-two bit registers, as shown in FIG. 17, in which a specific bit in each register corresponds to a specific queue. The registers are labeled Q-Out_Ready 745, Q-In_Ready 750, Q-Empty 755 and Q-Full 760. If a particular bit is set in the Q-Out_Ready register 750, the queue corresponding to that bit contains information that is ready to be read, while the setting of the same bit in the Q-In_Ready 752 register means
25 that the queue is ready to be written. Similarly, a positive setting of a specific bit in the Q-Empty register 755 means that the queue corresponding to that bit is empty, while a positive setting of a particular bit in the Q-Full register 760 means that the queue corresponding to that bit is full. Thus Q-Out_Ready 745 contains bits zero 746 through thirty-one 748, including bits twenty-seven 752, twenty-eight 754, twenty-nine 756 and thirty 758. Q-
30 In_Ready 750 contains bits zero 762 through thirty-one 764, including bits twenty-seven 766, twenty-eight 768, twenty-nine 770 and thirty 772. Q-Empty 755 contains bits zero 774 through thirty-one 776, including bits twenty-seven 778, twenty-eight 780, twenty-nine 782 and thirty 784, and Q-full 760 contains bits zero 786 through thirty-one 788, including bits twenty-seven 790, twenty-eight 792, twenty-nine 794 and thirty 796.

WO 00/13091

PCT/US98/24943

Q-zero, corresponding to FIFO 700, is a free buffer queue, which holds a list of addresses for all available buffers. This queue is addressed when the microprocessor or other devices need a free buffer address, and so commonly includes appreciable DRAM 460. Thus a device needing a free buffer address would check with Q-zero to obtain that address. Q-

5 twenty-seven, corresponding to FIFO 702, is a receive buffer descriptor queue. After processing a received frame by the receive sequencer the sequencer looks to store a descriptor for the frame in Q-twenty-seven. If a location for such a descriptor is immediately available in SRAM, bit twenty-seven 766 of Q-In_Ready 750 will be set. If not, the sequencer must wait for the queue manager to initiate a DMA move from SRAM to DRAM, thereby freeing

10 space to store the receive descriptor.

Operation of the queue manager, which manages movement of queue entries between SRAM and the processor, the transmit and receive sequencers, and also between SRAM and DRAM, is shown in more detail in FIG. 18. Requests which utilize the queues include Processor Request 802, Transmit Sequencer Request 804, and Receive Sequencer Request

15 806. Other requests for the queues are DRAM to SRAM Request 808 and SRAM to DRAM Request 810, which operate on behalf of the queue manager in moving data back and forth between the DRAM and the SRAM head or tail of the queues. Determining which of these various requests will get to use the queue manager in the next cycle is handled by priority logic Arbiter 815. To enable high frequency operation the queue manager is pipelined, with

20 Register A 818 and Register B 820 providing temporary storage, while Status Register 822 maintains status until the next update. The queue manager reserves even cycles for DMA, receive and transmit sequencer requests and odd cycles for processor requests. Dual ported QRAM 825 stores variables regarding each of the queues, the variables for each queue including a Head Write Pointer, Head Read Pointer, Tail Write Pointer and Tail Read Pointer

25 corresponding to the queue's SRAM condition, and a Body Write Pointer and Body Read Pointer corresponding to the queue's DRAM condition and the queue's size.

After Arbiter 815 has selected the next operation to be performed, the variables of QRAM 825 are fetched and modified according to the selected operation by a QALU 828, and an SRAM Read Request 830 or an SRAM Write Request 840 may be generated. The

30 variables are updated and the updated status is stored in Status Register 822 as well as QRAM 825. The status is also fed to Arbiter 815 to signal that the operation previously requested has been fulfilled, inhibiting duplication of requests. The Status Register 822 updates the four queue registers Q-Out_Ready 745, Q-In_Ready 750, Q-Empty 755 and Q-Full 760 to reflect the new status of the queue that was accessed. Similarly updated are

WO 00/13091

PCT/US98/24943

SRAM Addresses 833, Body Write Request 835 and Body Read Requests 838, which are accessed via DMA to and from SRAM head and tails for that queue. Alternatively, various processes may wish to write to a queue, as shown by Q Write Data 844, which are selected by multiplexor 846, and pipelined to SRAM Write Request 840. The SRAM controller services
 5 the read and write requests by writing the tail or reading the head of the accessed queue and returning an acknowledge. In this manner the various queues are utilized and their status updated.

FIGs. 19A-C show a least-recently-used register 900 that is employed for choosing which contexts or CCBs to maintain in INIC cache memory. The INIC in this embodiment
 10 can cache up to sixteen CCBs in SRAM at a given time, and so when a new CCB is cached an old one must often be discarded, the discarded CCB usually chosen according to this register 900 to be the CCB that has been used least recently. In this embodiment, a hash table for up to two hundred fifty-six CCBs is also maintained in SRAM, while up to two hundred fifty-six full CCBs are held in DRAM. The least-recently-used register 900 contains sixteen
 15 four-bit blocks labeled R0-R15, each of which corresponds to an SRAM cache unit. Upon initialization, the blocks are numbered 0-15, with number 0 arbitrarily stored in the block representing the least recently used (LRU) cache unit and number 15 stored in the block representing the most recently used (MRU) cache unit. FIG. 19A shows the register 900 at an arbitrary time when the LRU block R0 holds the number 9 and the MRU block R15 holds
 20 the number 6.

When a different CCB than is currently being held in SRAM is to be cached, the LRU block R0 is read, which in FIG. 19A holds the number 9, and the new CCB is stored in the SRAM cache unit corresponding to number 9. Since the new CCB corresponding to number 9 is now the most recently used CCB, the number 9 is stored in the MRU block, as shown in
 25 FIG. 19B. The other numbers are all shifted one register block to the left, leaving the number 1 in the LRU block. The CCB that had previously been cached in the SRAM unit corresponding to number 9 has been moved to slower but more cost-effective DRAM.

FIG. 19C shows the result when the next CCB used had already been cached in SRAM. In this example, the CCB was cached in an SRAM unit corresponding to number 10, and so after employment of that CCB, number 10 is stored in the MRU block. Only those
 30 numbers which had previously been more recently used than number 10 (register blocks R9-R15) are shifted to the left, leaving the number 1 in the LRU block. In this manner the INIC maintains the most active CCBs in SRAM cache.

WO 00/13091

PCT/US98/24943

In some cases a CCB being used is one that is not desirable to hold in the limited cache memory. For example, it is preferable not to cache a CCB for a context that is known to be closing, so that other cached CCBs can remain in SRAM longer. In this case, the number representing the cache unit holding the decacheable CCB is stored in the LRU block R0 rather than the MRU block R15, so that the decacheable CCB will be replaced immediately upon employment of a new CCB that is cached in the SRAM unit corresponding to the number held in the LRU block R0. FIG. 19D shows the case for which number 8 (which had been in block R9 in FIG. 19C) corresponds to a CCB that will be used and then closed. In this case number 8 has been removed from block R9 and stored in the LRU block R0. All the numbers that had previously been stored to the left of block R9 (R1-R8) are then shifted one block to the right.

FIG. 20 shows some of the logical units employed to operate the least-recently-used register 900. An array of sixteen, three or four input multiplexors 910, of which only multiplexors MUX0, MUX7, MUX8, MUX9 and MUX15 are shown for clarity, have outputs fed into the corresponding sixteen blocks of least-recently-used register 900. For example, the output of MUX0 is stored in block R0, the output of MUX7 is stored in block R7, etc. The value of each of the register blocks is connected to an input for its corresponding multiplexor and also into inputs for both adjacent multiplexors, for use in shifting the block numbers. For instance, the number stored in R8 is fed into inputs for MUX7, MUX8 and MUX9. MUX0 and MUX15 each have only one adjacent block, and the extra input for those multiplexors is used for the selection of LRU and MRU blocks, respectively. MUX15 is shown as a four-input multiplexor, with input 915 providing the number stored on R0.

An array of sixteen comparators 920 each receives the value stored in the corresponding block of the least-recently-used register 900. Each comparator also receives a signal from processor 470 along line 935 so that the register block having a number matching that sent by processor 470 outputs true to logic circuits 930 while the other fifteen comparators output false. Logic circuits 930 control a pair of select lines leading to each of the multiplexors, for selecting inputs to the multiplexors and therefore controlling shifting of the register block numbers. Thus select lines 939 control MUX0, select lines 944 control MUX7, select lines 949 control MUX8, select lines 954 control MUX9 and select lines 959 control MUX15.

When a CCB is to be used, processor 470 checks to see whether the CCB matches a CCB currently held in one of the sixteen cache units. If a match is found, the processor sends a signal along line 935 with the block number corresponding to that cache unit, for example

WO 00/13091

PCT/US98/24943

number 12. Comparators 920 compare the signal from that line 935 with the block numbers and comparator C8 provides a true output for the block R8 that matches the signal, while all the other comparators output false. Logic circuits 930, under control from the processor 470, use select lines 959 to choose the input from line 935 for MUX15, storing the number 12 in
 5 the MRU block R15. Logic circuits 930 also send signals along the pairs of select lines for MUX8 and higher multiplexors, aside from MUX15, to shift their output one block to the left, by selecting as inputs to each multiplexor MUX8 and higher the value that had been stored in register blocks one block to the right (R9-R15). The outputs of multiplexors that are to the left of MUX8 are selected to be constant.

10 If processor 470 does not find a match for the CCB among the sixteen cache units, on the other hand, the processor reads from LRU block R0 along line 966 to identify the cache corresponding to the LRU block, and writes the data stored in that cache to DRAM. The number that was stored in R0, in this case number 3, is chosen by select lines 959 as input 915 to MUX15 for storage in MRU block R15. The other fifteen multiplexors output to their
 15 respective register blocks the numbers that had been stored each register block immediately to the right.

For the situation in which the processor wishes to remove a CCB from the cache after use, the LRU block R0 rather than the MRU block R15 is selected for placement of the number corresponding to the cache unit holding that CCB. The number corresponding to the
 20 CCB to be placed in the LRU block R0 for removal from SRAM (for example number 1, held in block R9) is sent by processor 470 along line 935, which is matched by comparator C9. The processor instructs logic circuits 930 to input the number 1 to R0, by selecting with lines 939 input 935 to MUX0. Select lines 954 to MUX9 choose as input the number held in register block R8, so that the number from R8 is stored in R9. The numbers held by the other
 25 register blocks between R0 and R9 are similarly shifted to the right, whereas the numbers in register blocks to the right of R9 are left constant. This frees scarce cache memory from maintaining closed CCBs for many cycles while their identifying numbers move through register blocks from the MRU to the LRU blocks.

All told, the above-described devices and systems for processing of data
 30 communication result in dramatic reductions in the time required for processing large, connection-based messages. Protocol processing speed is tremendously accelerated by specially designed protocol processing hardware as compared with a general purpose CPU running conventional protocol software, and interrupts to the host CPU are also substantially reduced. These advantages can be provided to an existing host by addition of an intelligent

WO 00/13091

PCT/US98/24943

network interface card (NIC), or the protocol processing hardware may be integrated with the CPU. In either case, the protocol processing hardware and CPU intelligently decide which device processes a given message, and can change the allocation of that processing based upon conditions of the message.

WO 00/13091

PCT/US98/24943

Claims

1. A method for communication between a network and a host computer having a processor and a sequential stack of protocol layers, the method comprising:
receiving, by said host from said network, a message packet including data and a plurality of headers corresponding to said stack of protocol layers, said data intended for placement in a destination of said host according to protocol processing of said headers, processing, as a group and at one time, said plurality of headers, including creating a summary of said group of headers, choosing, based upon said summary, whether to process said packet by said protocol layers, and sending said data to said destination according to said summary of said group of headers, whereby sequential processing of said packet by said stack of protocol layers is avoided.
2. The method of claim 1, wherein said processing of said group of headers occurs during said receiving, by said host from said network, of said message packet.
3. The method of claim 1, further comprising creating a communication control block for a connection including said packet, and matching said summary with said communication control block, for sending said data to said destination.
4. The method of claim 1, further comprising creating a communication control block for a connection including said packet, wherein sending said data to said destination includes guiding said data by said communication control block.
5. The method of claim 8, further comprising transmitting a second message packet from said host to said network by referencing said communication control block.

WO 00/13091

PCT/US98/24943

6. A method for processing communication between a network and a host having a sequential protocol processing stack, the method comprising:

providing a device including a communication processor, said device being connected to said host and said network,

receiving a message frame from said network by said host, said frame including data and a series of headers corresponding to said sequential protocol processing stack,

analyzing said series of headers as a stream of bytes by said device, including processing said headers without copying said data, thereby creating a summary of said frame, and

selecting, based upon said processing, whether to process said packet by said stack or to send said data to a destination according to said summary.

7. The method of claim 6, further comprising:

creating, by said host, a communication control block for a message including said frame,

storing said communication control block in said device, and

guiding said data to a destination denoted by said communication control block.

8. The method of claim 7, further comprising comparing said summary with said communication control block, prior to guiding said data to said destination.

9. The method of claim 6, further comprising:

transmitting, via said device, transmission data from said host to said network, including simultaneously prepending several protocol headers to said transmission data for network transfer to a remote host.

WO 00/13091

PCT/US98/24943

10. A method for communication between a network and a host computer having a processor and a sequential stack of protocol layers, the method comprising:

receiving, by said host from said network, a message having multiple packets, each of said packets including a data portion and an associated sequence of headers which include information corresponding to said sequential stack of protocol layers and indicate an upper layer destination in said host for said data, and

sending a plurality of said data portions to said destination without said associated headers and without generating an interrupt to any host CPU.

11. The method of claim 10, further comprising choosing whether to process said packets by said stack of protocol layers, prior to sending said data portions to said destination.

12. The method of claim 10, further comprising providing to said host a protocol processing device, and

summarizing said headers with said device, prior to sending said data portions to said destination without said headers.

13. The method of claim 10, further comprising transmitting a data file from said host to said network, including dividing said data file into a series of data units, prepending headers to said data units and thereby creating a series of network frames, and placing said network frames on said network without generating an interrupt to any host CPU.

WO 00/13091

PCT/US98/24943

14. A method for communication between a host computer and a network, the host computer having a CPU, a storage unit and a sequential stack of protocol layers, the method comprising:

- providing a device connected to said network and said host, said device having a processor,
- receiving by said device a first message from said network,
- processing said first message, including creating a communication control block for said first message,
- receiving by said device a second message from said network, said second message including data and a header, said header including a series of protocol layer headers,
- processing said header by said device, including generating a summary of said header, without copying said data during said processing of said header, and
- sending said data by said device to an upper layer of said protocol layers in a form suitable for said upper layer, including guiding said sending with said communication control block.

15. The method of claim 14, further comprising receiving by said device a third message relating to said first and second messages, and passing said communication control block from said device to said storage unit, thereby passing control of processing said third message to said CPU.

16. The method of claim 14, further comprising matching said summary with said communication control block, prior to sending said data to said upper layer.

17. The method of claim 14, further comprising transmitting from said host to said network a third message, including sending said third message via said device by referencing said communication control block and prepending a transmission header to data acquired from a host source, said transmission header including a plurality of protocol layer headers.

WO 00/13091

PCT/US98/24943

18. A method for communication between a local host and a remote host connected by a network, with the local host having a protocol processing stack and an associated protocol processing device, the method comprising:

creating, by the protocol processing stack, a communication control block defining a connection between the local host and the remote host,

passing said communication control block to the device, and thereby passing control of processing a message packet associated with said connection and transferred between the network and the local host, such that said packet is processed by the device instead of by the protocol processing stack.

19. The method of claim 18, further comprising passing said communication control block back to the local host, such that a second message packet transferred between the network and the local host and associated with said connection is generally processed by the protocol processing stack.

20. The method of claim 18, further comprising:

receiving, by the device, a message frame from the network, and

summarizing, by the device, said message frame, thereby generating a summary of said message frame, and

comparing said summary with said communication control block.

21. The method of claim 18, further comprising:

transmitting, by the device, a message frame to the network, including forming a header based upon said communication control block and prepending said header to said message frame.

WO 00/13091

PCT/US98/24943

22. A method for network communication by a host computer having a processor, a memory and a sequential stack of protocol layers, the method comprising:

receiving by the host from the network a packet including data and a plurality of headers relating to the stack of protocol layers, said data having a destination in said host,
categorizing said packet with a hardware logic sequencer, including
classifying said headers and creating a summary of said packet, and
choosing, based upon said summary, whether to send said packet to said stack of protocol layers or to bypass said stack of protocol layers by sending said data to said destination.

23. The method of claim 22, further comprising:

sending said packet to said stack of protocol layers,
processing said packet with said stack of protocol layers and thereby creating a context for said message,
receiving by said host from said network a related packet including additional data and additional headers, and
employing said context for sending said related packet to said destination without processing said packet by said stack of protocol layers.

24. The method of claim 22, further comprising creating a context for a message including said packet, said context defining a connection between said host and a remote host, wherein choosing whether to send said packet to said stack of protocol layers or to bypass said stack of protocol layers includes comparing said summary with said context.

25. The method of claim 22, further comprising bypassing said stack of protocol layers by sending said data to said destination in a form suitable for said destination.

WO 00/13091

PCT/US98/24943

26. A device for use with a local host that is connectable to a remote host via a network, the local host containing a CPU operating a stack of protocol processing layers that define a connection context between an application of the local host and an application of the remote host, the device comprising:

a communication processing mechanism connected to the CPU and to the network and containing a processor configured for choosing, by referencing the context, whether to process a network message by the protocol processing layers or to avoid the protocol processing layers and employ the context for transferring data contained in said message between the network and the local host application.

27. The device of claim 26, wherein said communication processing mechanism includes a receive sequencer connected to said processor and configured for validating a message packet received from the network, and said message packet contains control information corresponding to the stack of protocol layers.

28. The device of claim 26, wherein said communication processing mechanism includes a receive sequencer connected to said processor and configured for generating a summary of a message packet received from the network, said message packet containing control information corresponding to the stack of protocol layers, with said processor adapted for comparing said summary with said context.

29. The device of claim 26, wherein said processor is adapted for creating a header corresponding to said context and including control information corresponding to several of the protocol processing layers, and prepending said header to said data for transmission of said message from the host to the remote host.

30. The device of claim 26, wherein said communication processing mechanism has a direct memory access unit for sending, based upon said context, said data from said communication processing mechanism to the host application, without a header accompanying said data.

WO 00/13091

PCT/US98/24943

31. The device of claim 27, wherein said processor includes a plurality of microprocessors, with at least one of said microprocessors primarily adapted for processing messages received by the host from the network and a second of said microprocessors adapted for processing messages transmitted from the host to the network.

32. The device of claim 31, wherein said microprocessors utilize shared hardware functions in rotating phases.

33. The device of claim 26, wherein said communication processing mechanism includes a queue manager configured for queuing information in a plurality of queue storage units, wherein at least one of said queue storage units contains SRAM and DRAM.

34. A communication device for a host computer connectable to a network, the host computer having a CPU with a stack of protocol processing layers operable by the CPU for processing network communications and defining a connection context between a destination in the host and a source in a remote host, the device comprising:

a receive sequencer connected to the network and configured for validating a message packet received from the network by the device, said message packet containing data and a header with control information regarding several of said protocol layers, said receive sequencer adapted for creating a summary of said packet for storage in the device, and

a communication processor connected to said receive sequencer and to the host, said communication processor adapted for comparing said summary with the communication control block and choosing, based upon said comparing, whether to send said message packet to the stack for protocol processing or to send said data directly to the destination without processing said message packet by the stack.

35. The device of claim 34, wherein said communication processor contains a plurality of pipelined microprocessors operating in rotating phases, with at least one of said microprocessors configured for processing messages received by the host from the network and at least one of said microprocessors configured for processing messages transmitted from the host to the network.

WO 00/13091

PCT/US98/24943

36. The device of claim 9, further comprising a queue manager containing logic circuits configured for queuing information for said communication processor and said receive sequencer in a plurality of storage queues.
37. The device of claim 36, wherein said queues contain both SRAM and DRAM.
38. The device of claim 34, further comprising a cache memory and a plurality of connection contexts, with a hash of each said connection context being stored in said cache memory for comparison with said summary.
39. The device of claim 38, further comprising a least recently used register for determining which of said connection contexts are stored in said cache memory.
40. The device of claim 34, wherein said receive sequencer has a sequence of hardware logic units for processing a header contained in said message packet and having control information substantially corresponding to said stack of protocol processing layers.
41. The device of claim 34, wherein said communication processor commands a direct memory access controller for sending, based upon said context, said data to the host destination without said header.
42. The device of claim 34, wherein said communication processor is configured for transmitting a second message and thereby transferring transmit data from the host destination to the network, including prepending a header derived from said context to said transmit data.

WO 00/13091

PCT/US98/24943

43. A device for communication between a local host with a CPU and a remote host, the hosts connected by a network, the device comprising:

a communication processing mechanism connected to the network and to the local host, said mechanism including hardware logic for processing a data packet and generating a summary of said packet,

a protocol processing stack disposed in the local host and operable by the CPU for creating a communication control block and passing said communication control block to said mechanism, with said communication control block defining a connection between the local host and the remote host,

wherein said mechanism and said protocol processing stack are arranged such that a message corresponding to said connection and transferred between said network and said local host is processed by said mechanism instead of by the CPU when said mechanism is holding said communication control block.

44. The device of claim 43, wherein said mechanism has a plurality of network connections.

45. A device for processing communication between a network and a host having a stack of protocol layers, said device comprising:

a plurality of logic units for categorizing a message packet received from the network, said packet including data and a header and flowing through said logic units as a stream of bits, with said logic units creating a summary of said packet from said stream,

a memory for storing said packet and said summary, and

a microprocessor for matching said summary with a connection context, and for moving said packet without said protocol information to a destination in the host indicated by said context.

46. The device of claim 45, wherein said microprocessor includes a plurality of pipelined processors, with one of said processors configured for transmitting network messages and another of said processors configured for receiving network messages.

47. The device of claim 45, further comprising a CPU operating a stack of protocol processing layers for processing a second packet having a second summary not matching said connection context.

WO 00/13091

PCT/US98/24943

48. A device for transferring a message between a network and a host computer, the message including a series of headers corresponding to a sequence of protocol layers, the device comprising a series of sequencers aligned for processing the series of headers as a stream of bits and generating a status of said message, with at least one of said sequencers including logic for categorizing a header corresponding to an upper layer of said protocol layers.
49. The device of claim 48, further comprising a communication processor connected to said sequencers and capable of employing said status for sending said message without the headers to a destination of said host.
50. A device for transmitting messages between a network and a host, the device comprising:
- an array of variable length FIFO circuits defining a plurality of queues,
 - a receive sequencer configured for generating a status of a frame of the messages received from the network and storing said status in at least one of said queues, and
 - a protocol processor containing a plurality of pipelined microprocessors operating a set of logical units in alternating phases, with a first of said microprocessors adapted for processing the messages received from the network and a second of said microprocessors adapted for processing the messages transmitted to the network, said processor employing a communication control block in concert with said sequencer for accelerating transfer of the messages.
51. The device of claim 50, wherein at least one of said queues includes a DRAM and an SRAM storage unit.

WO 00/13091

PCT/US98/24943

52. A device for communication between a network and a host computer having a processor and a sequential stack of protocol layers, the device comprising:

means for receiving, by said host from said network, a message packet including data and a plurality of headers corresponding to said stack of protocol layers, said data intended for placement in a destination of said host according to protocol processing of said headers,

means for processing said plurality of headers, including creating a summary of said group of headers, and

means for sending said data to said destination according to said summary of said group of headers, whereby sequential processing of said packet by said stack of protocol layers is avoided.

1/18




FIG. 1




FIG. 2

2/18

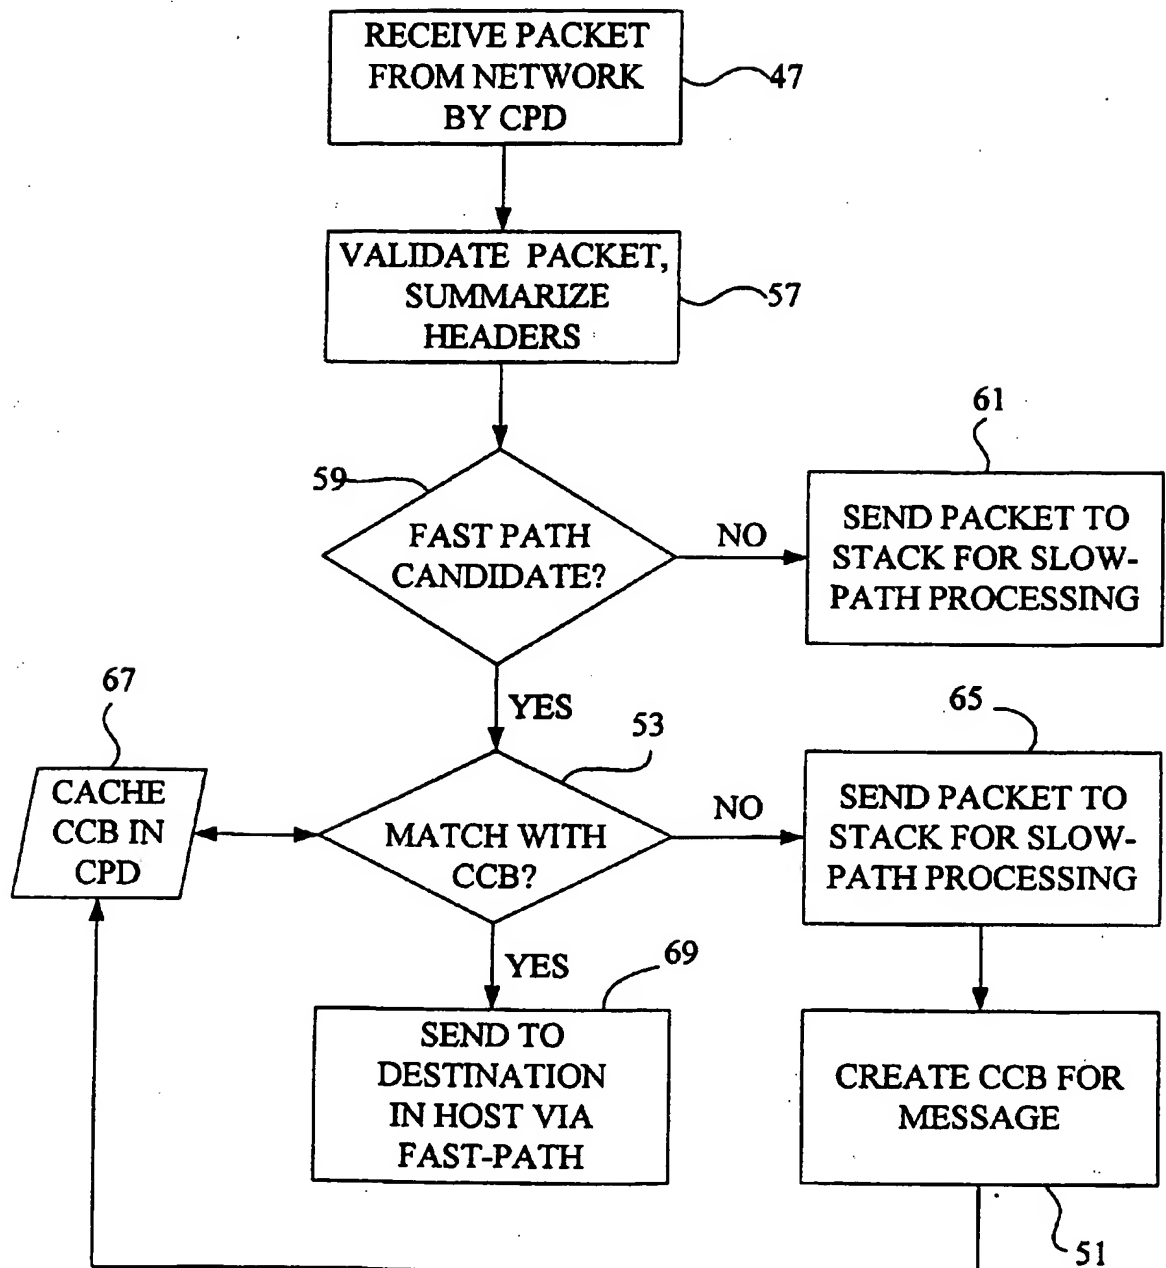


FIG. 3

3/18




FIG. 4A




FIG. 4B

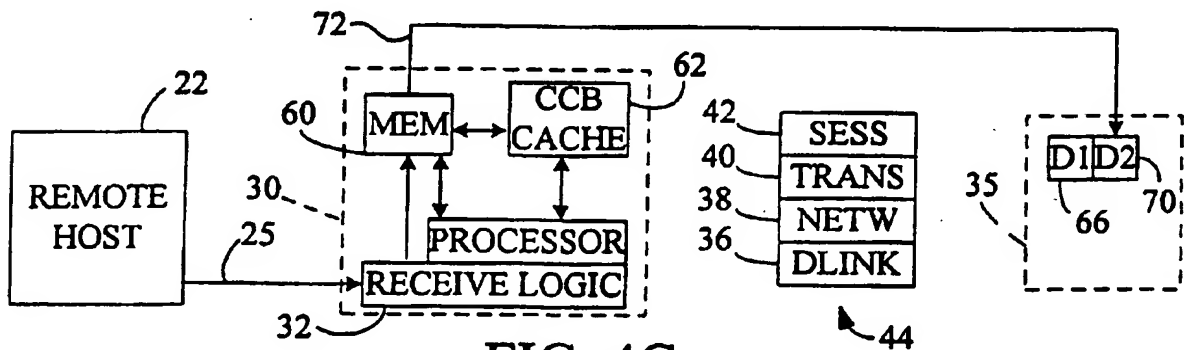


FIG. 4C




FIG. 4D

4/18




FIG. 5




FIG. 6

5/18




FIG. 7

6/18




FIG. 8

7/18




FIG. 9

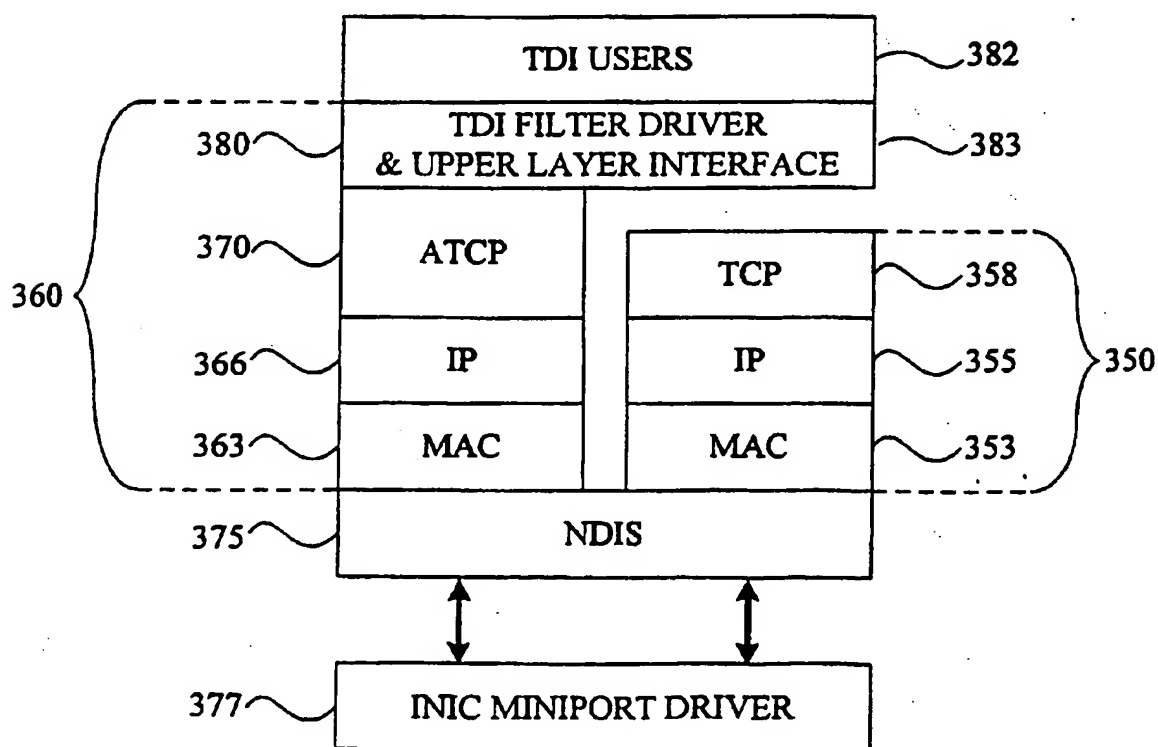


FIG. 11

8/18

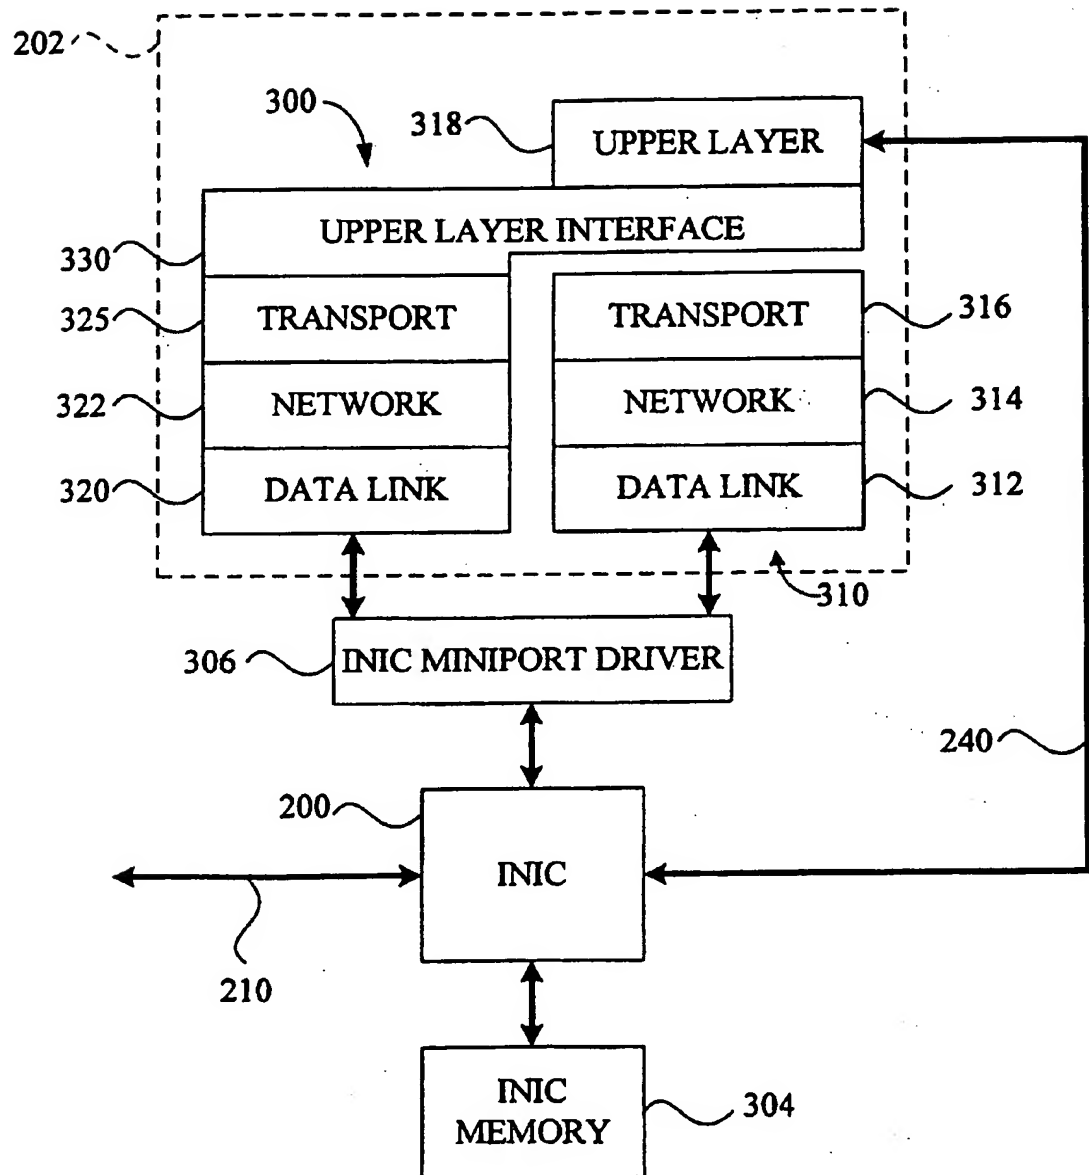


FIG. 10

9/18




FIG. 12

10/18

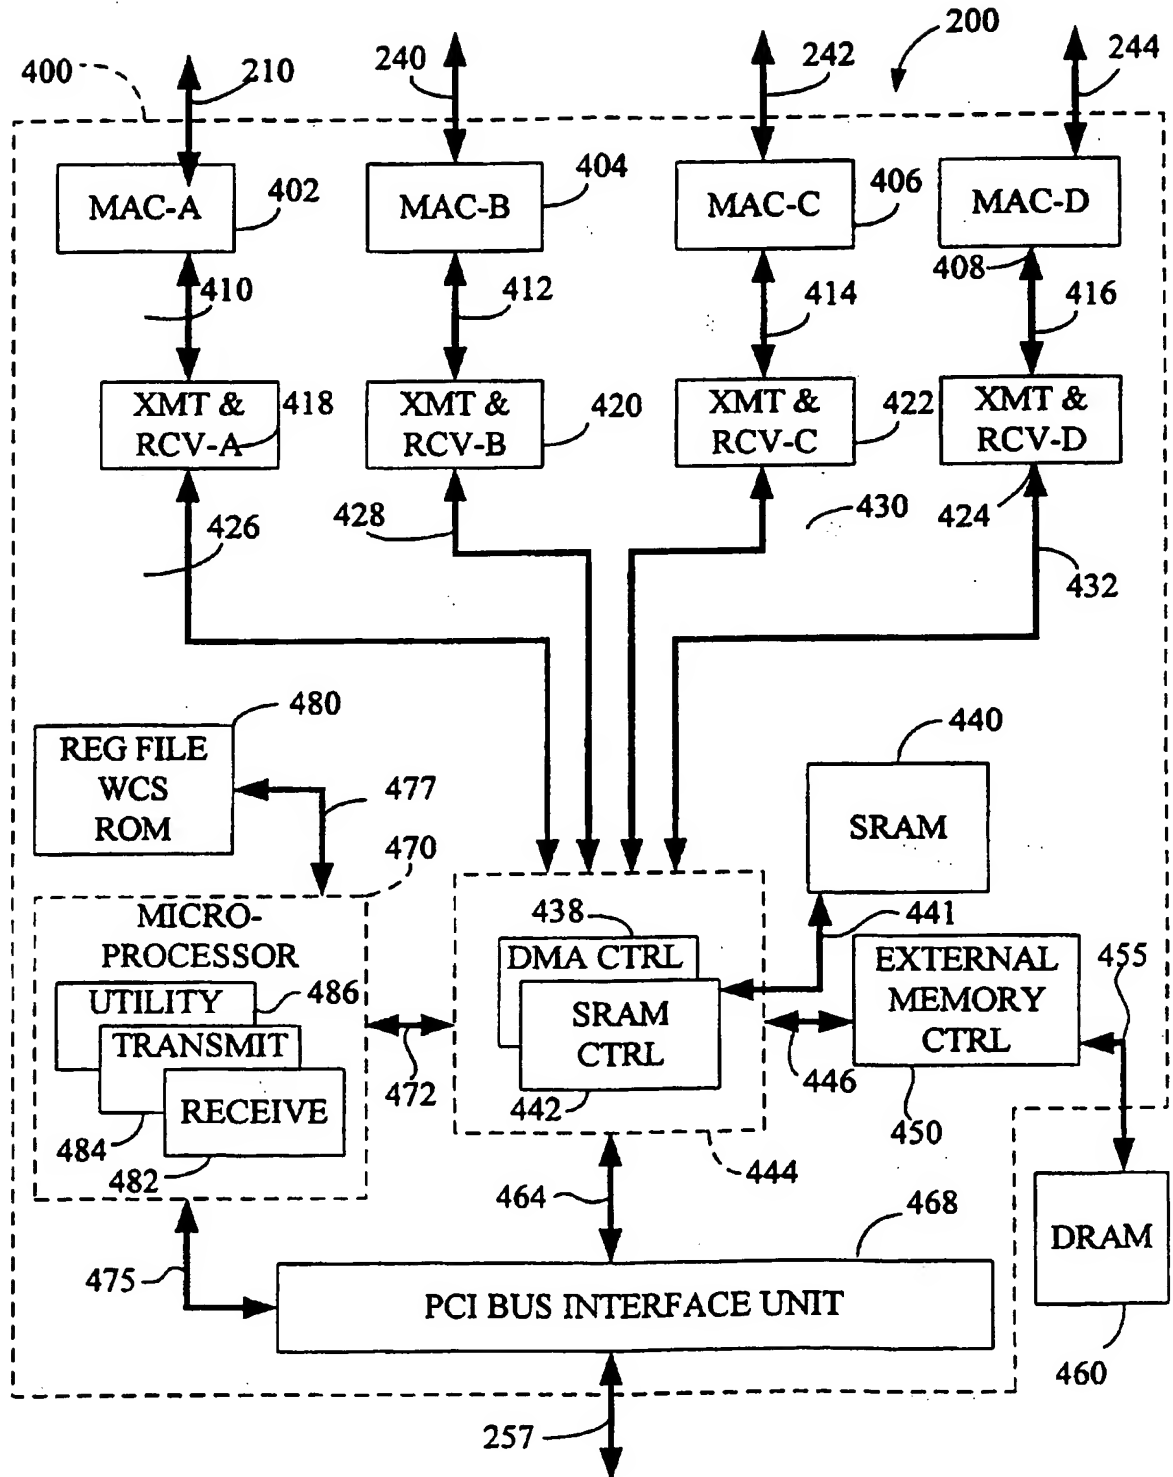


FIG. 13

11/18




FIG. 14

12/18

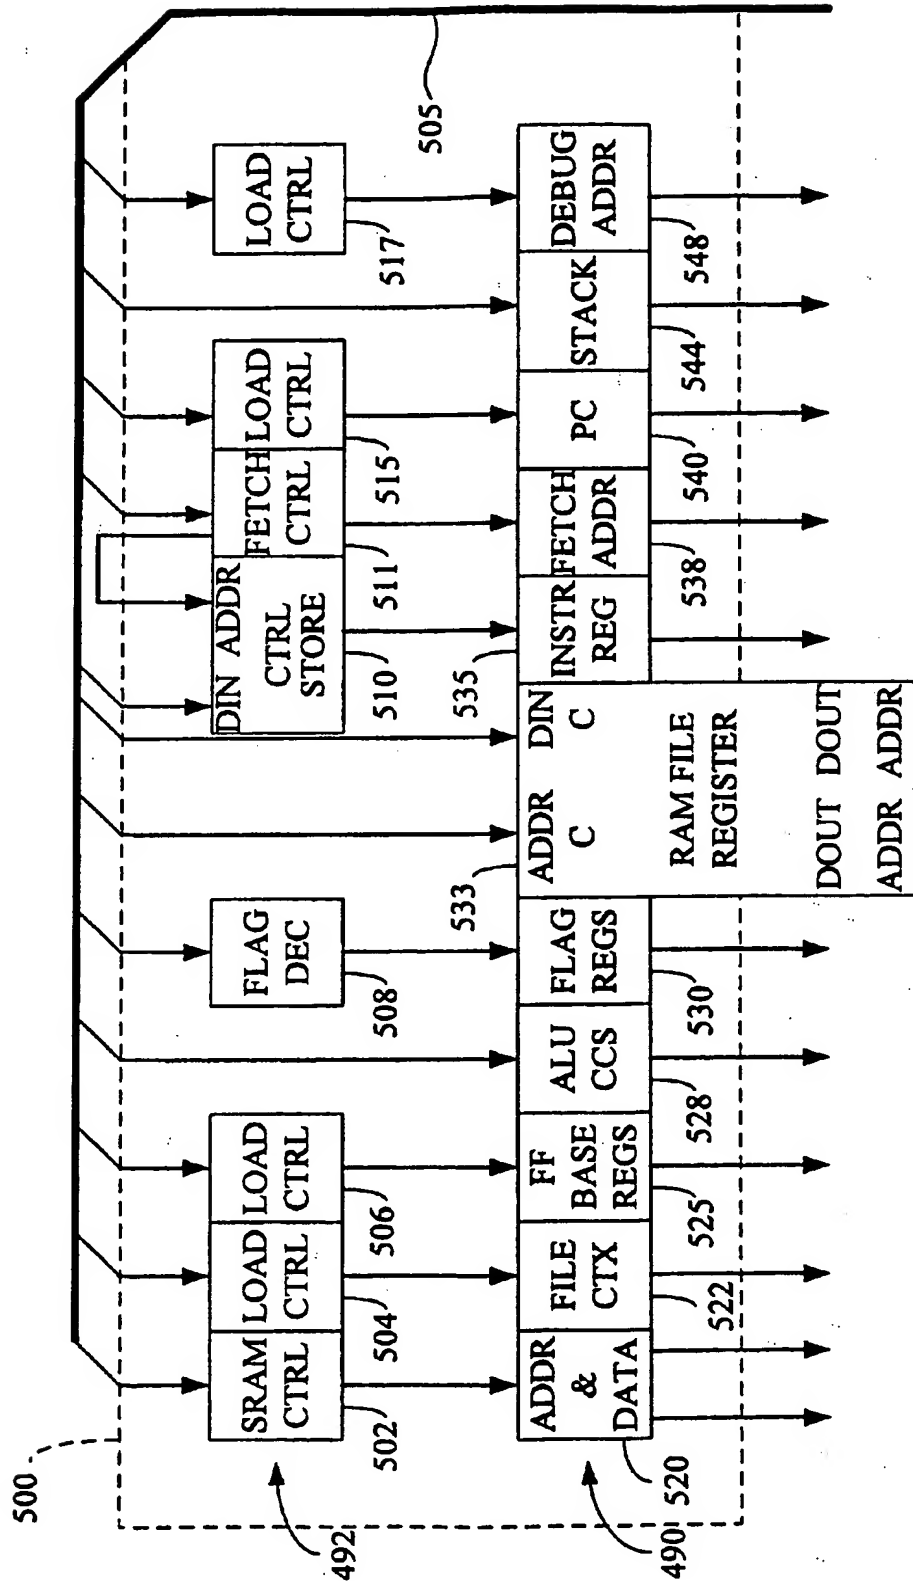


FIG. 15A

13/18

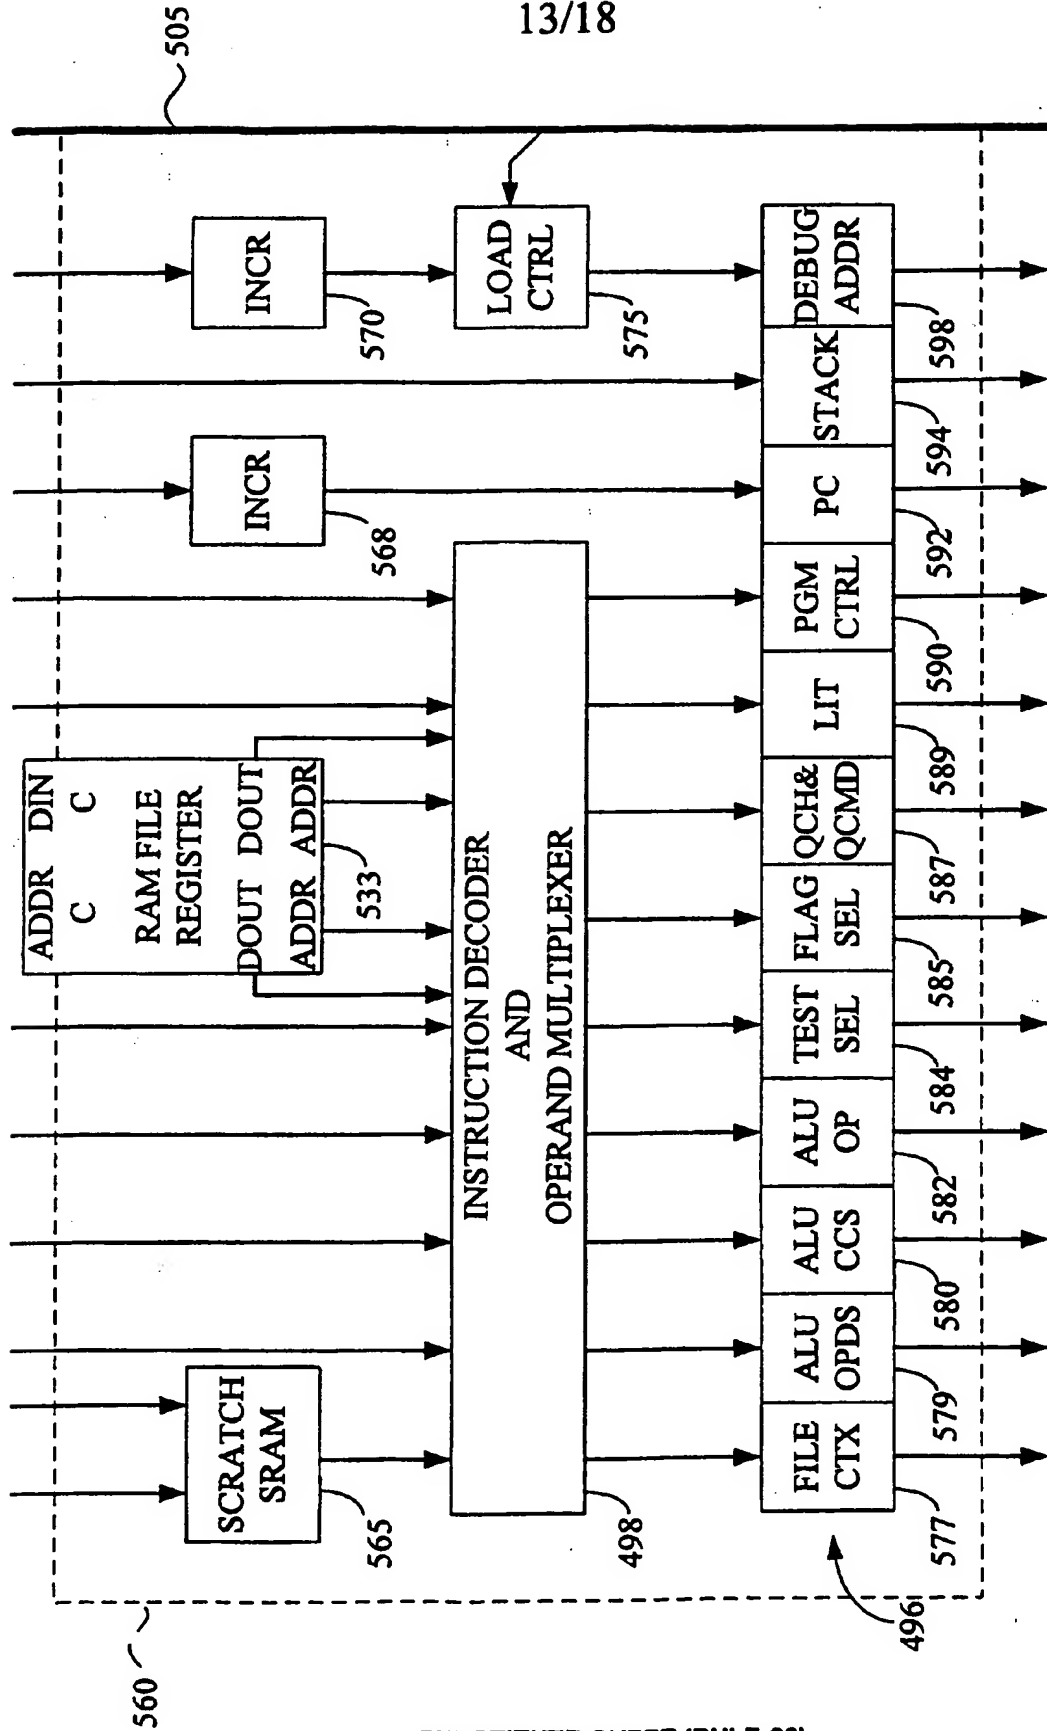


FIG. 15B

14/18

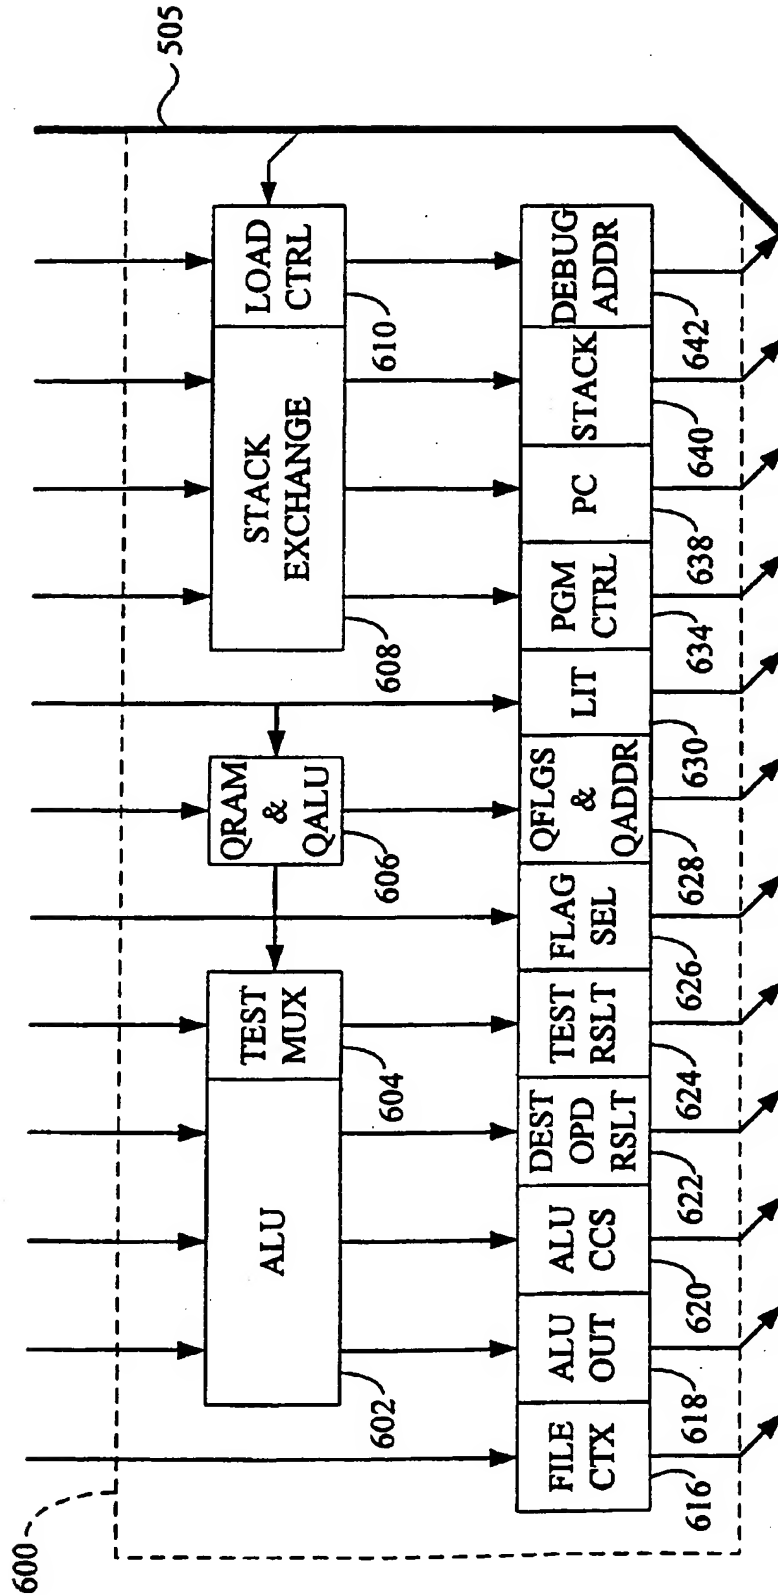


FIG. 15C

15/18




FIG. 16




FIG. 17

16/18








FIG. 18

17/18






FIG. 20